JMLR: Workshop and Conference Proceedings vol 65:1–23, 2017

Stochastic Approximation for Canonical Correlation Analysis

Raman Arora ARORA@CS.JHU.EDU
Teodor V. Marinov TMARINO2@JHU.EDU
Poorya Mianjy MIANJY@JHU.EDU
Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218

arXiv:1702.06818v1 [cs.LG] 22 Feb 2017

Abstract

We study canonical correlation analysis (CCA) as a stochastic optimization problem. We show that regularized CCA is efficiently PAC-learnable. We give stochastic approximation (SA) algorithms that are instances of stochastic mirror descent, which achieve ϵ-suboptimality in the population objective in time $\mathrm{poly}(\frac{1}{\epsilon}, \frac{1}{\delta}, d)$ with probability $1-\delta$, where d is the input dimensionality.

Keywords: stochastic approximation, canonical correlation analysis, representation Learning

1. Introduction

Canonical Correlation Analysis (CCA) (Hotelling, 1936) is a ubiquitous statistical technique for finding linear components of two sets of random variables that are maximally correlated. CCA is often posed as a dimensionality reduction problem about a fixed dataset of n paired data points. In this paper, we take a stochastic optimization view of CCA, rather than thinking of it as a dimensionality reduction technique for a given finite data set. Following Arora et al. (2012), we pose CCA as the following stochastic optimization problem: given a pair of random vectors $(\mathrm{x}, \mathrm{y}) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y}$, with some (unknown) joint distribution $\mathscr{D}$, find matrices $\mathrm{U} \in \mathbb{R}^{d_x \times k}$ and $\mathrm{V} \in \mathbb{R}^{d_y \times k}$ that solve

$$\begin{aligned} \underset{\mathrm{U}\in\mathbb{R}^{d_x\times k},\, \mathrm{V}\in\mathbb{R}^{d_y\times k}}{\text{maximize}} \quad & \mathbb{E}_{\mathrm{x},\mathrm{y}}[\mathrm{x}^\top \mathrm{U}\mathrm{V}^\top \mathrm{y}] \\ \text{subject to} \quad & \mathrm{U}^\top \mathbb{E}_{\mathrm{x}}[\mathrm{x}\mathrm{x}^\top]\mathrm{U} = \mathrm{V}^\top \mathbb{E}_{\mathrm{y}}[\mathrm{y}\mathrm{y}^\top]\mathrm{V} = \mathrm{I}_k \end{aligned} . \tag{1}$$

The advantage of viewing CCA as a stochastic optimization problem is two-fold. First, from a learning perspective, we argue that unlike typical unsupervised learning problems where notion of learnability is poorly understood, we can describe generalization error, and therefore define learnability, in terms of ϵ-suboptimality of the solution to stochastic CCA problem in equation (1). Second, from a computational perspective, we argue that in a big data setting, where we have as much data as we would like, and the goal is to find a "good" enough subspace in terms of capturing correlation rather than a "true" subspace, the stochastic optimization view motivates stochastic approximation algorithms that can easily scale to very large datasets. Roughly speaking, a Stochastic Approximation (SA) algorithm is an iterative algorithm, where in each iteration a single sample from the population is used to perform an update, as in stochastic gradient descent (SGD), the classic SA algorithm.

There are several computational challenges associated with solving Problem 1 when viewed as a stochastic optimization problem. A first challenge stems from the fact that Problem 1 is non-convex whereas most learning problems that we can learn efficiently are convex; in fact, further constraints such as smoothness or Lipshcitz-boundedness are necessary even for efficiently learning convex

© 2017 R. Arora, T.V. Marinov & P. Mianjy.

problems (Shalev-Shwartz and Ben-David, 2014). Fortunately, like several related problems such as principal component analysis (PCA), the solution to Problem 1 can be given as a generalized eigenvalue problem. Numerical techniques based on power iteration method and its variants can be applied to these problems to find globally optimal solutions. Therefore, much of the recent work has focused on analyzing optimization error for power iteration method for the generalized eigenvalue problem (Ge et al., 2016; Wang et al., 2016; Allen-Zhu and Li, 2016). Note, however, that these works give guarantees on *numerical optimization error* for finding left and right singular vectors of a fixed given matrix based on empirical estimates of the covariance matrices. It is not immediately clear how one would give generalization error bounds for Problem 1 using existing techniques.

A second computational challenge is somewhat unique to the CCA optimization problem. The main difficulty here, compared to PCA, and most other learning problems, is that the constraints also involve stochastic quantities and depend on the unknown distribution $\mathscr{D}$. An alternative formulation would be to normalize the objective with product of the standard deviations, instead of including the uncorrelated-ness constraints; e.g., for $k = 1$, we can consider the objective $\rho(\mathrm{u}^T\mathrm{x}, \mathrm{v}^T\mathrm{y}) = \mathbb{E}_{\mathrm{x},\mathrm{y}}\left[\mathrm{u}^\top \mathrm{x}\mathrm{y}^\top \mathrm{v}\right] / \left(\sqrt{\mathbb{E}_{\mathrm{x}}\left[\mathrm{u}^\top \mathrm{x}\mathrm{x}^\top \mathrm{u}\right]}\sqrt{\mathbb{E}_{\mathrm{y}}\left[\mathrm{v}^\top \mathrm{y}\mathrm{y}^\top \mathrm{v}\right]}\right)$. This yields an unconstrained optimization problem. However, the objective is no longer an expectation, but is instead a ratio of expectations. This departs significantly from the typical stochastic approximation scenario. Crucially, it is not at all clear how to obtain unbiased estimates of the gradient of the objective $\rho(\mathrm{u}^T\mathrm{x}, \mathrm{v}^T\mathrm{y})$. A possible, though not completely satisfying, approach might be to consider pre-whitened data in each view (or add a normalization pre-processing step). In fact, in many applications of CCA to genomics and NLP, it is often assumed that data is pre-normalized (Witten et al., 2009; Lykou and Whittaker, 2010; Balakrishnan et al., 2012; Cohen et al., 2012). Under such an assumption the CCA problem reduces to partial least squares (PLS) and we can use SA approaches discussed in (Arora et al., 2016).

Finally, it can be shown that the CCA problem given in Problem 1 is non-learnable, i.e. there exists distributions on which we may incur small empirical error on the training sample but arbitrary large generalization error. This is not surprising since linear regression even in one-dimensional setting is known to be non-learnable unless we consider some form of regularization, e.g. ℓ_2 regularization as in ridge regression. Similarly, we consider the regularized CCA problem (see Section 2), unlike previous work which focus on the numerical optimization view of the problem rather than stochastic optimization (Ge et al., 2016; Allen-Zhu and Li, 2016).

1.1. Main contributions

The main contributions of the paper are as follows.

- Previous works on CCA can be classified as instances of Sample Average Approximation (SAA), where the focus is on optimizing the empirical objective. The runtime for these approaches scales poorly with sample size making them intractable for large datasets. We study CCA as a stochastic optimization problem. The goal here is to optimize a "population objective" based on i.i.d. draws from the population rather than capturing the *sample*, i.e. the training error. This justifies stochastic approximation approaches that are far from optimal on the sample but are essentially as good as the SAA approach on the population. Such a view has been advocated in supervised machine learning (Bousquet and Bottou, 2008; Shalev-Shwartz and Srebro, 2008); here, we carry over the same view to the rich world of unsupervised learning.

- We give a convex relaxation of the CCA optimization problem. Furthermore, we present two stochastic approximation algorithms for solving the resulting problem – these algorithms are instances of stochastic mirror descent with the choice of potential function being Frobenius norm and von Neumann entropy, respectively. We emphasize that these algorithms work in a streaming setting, i.e. process one sample at a time, requiring only a single pass through the data, and can easily scale to large datasets.

- We give precise sample complexity and iteration complexity bounds for our algorithms. In other words, we give upper bounds on samples needed and overall runtime needed to guarantee a user-specified generalization error for CCA. These bounds are gap-free and to the best of our knowledge first characterization of CCA as a stochastic optimization problem.

- We analyze the generalization error incurred by the current state of the art methods for solving the empirical problem (Wang et al., 2016; Ge et al., 2016; Allen-Zhu and Li, 2016). In particular, we show that the error decays as $\mathcal{O}\left(\frac{1}{\text{gap}\sqrt{t}}\right)$, where t is the number of samples.

- We show empirically that the proposed algorithms compare well and in certain cases outperform existing state-of-the-art methods for stochastic optimization of CCA on both synthetic and real datasets. We will make our implementation of the proposed algorithms and existing techniques available online[1].

1.2. Notation

Scalars, vectors and matrices are represented by normal, Roman and capital Roman letters respectively, e.g. x, x, and X. I_k denotes identity matrix of size k, where we drop the subscript whenever the size is clear from the context. ℓ^2-norm of a vector x is denoted by $\|\mathrm{x}\|$. For any matrix X, spectral norm, nuclear norm, and Frobenius norm are represented by $\|\mathrm{X}\|_2$, $\|\mathrm{X}\|_F$, and $\|\mathrm{X}\|_*$ respectively. Trace of a square matrix X is denoted by $\mathrm{Tr}\left(\mathrm{X}\right)$. Given two matrices $\mathrm{X} \in \mathbb{R}^{k\times d}, \mathrm{Y} \in \mathbb{R}^{k\times d}$, the standard inner-product between the two is given as $\langle \mathrm{X}, \mathrm{Y}\rangle = \mathrm{Tr}\left(\mathrm{X}^\top \mathrm{Y}\right)$. For self-adjoint X and Y, we say $\mathrm{X} \succeq \mathrm{Y}$ if $\mathrm{X} - \mathrm{Y}$ is positive semi-definite (PSD). Let singular value decomposition (SVD) of arbitrary X be given as $\mathrm{X} = \mathrm{U}\Sigma\mathrm{V}^\top$. Then $\mathbb{R} \to \mathbb{R}$ functions such as power, logarithm and exponential are extended to X by acting point-wise on diagonal elements of Σ. Let $\mathrm{x} \in \mathbb{R}^{d_x}$ and $\mathrm{y} \in \mathbb{R}^{d_y}$ denote two sets of centered random variables jointly distributed as $\mathscr{D}$ with corresponding auto-covariance matrices $\mathrm{C}_{xx} = \mathbb{E}_{\mathrm{x}}[\mathrm{x}\mathrm{x}^\top]$, $\mathrm{C}_{yy} = \mathbb{E}_{\mathrm{y}}[\mathrm{y}\mathrm{y}^\top]$, and cross-covariance matrix $\mathrm{C}_{xy} = \mathbb{E}_{(\mathrm{x},\mathrm{y})}[\mathrm{x}\mathrm{y}^\top]$. We denote the regularized auto-covariance matrices by $\mathrm{C}_x = \mathrm{C}_{xx} + r_x\mathrm{I}$ and $\mathrm{C}_y = \mathrm{C}_{yy} + r_y\mathrm{I}$, where $r_x, r_y \in \mathbb{R}_+$ are regularization parameters. We denote the empirical estimates from t points of the auto and cross-covariance operators as $\mathrm{C}_{xx,t}, \mathrm{C}_{yy,t}$ and $\mathrm{C}_{xy,t}$. The regularized versions of the empirical auto-covariance operators are denoted as $\mathrm{C}_{x,t}$ and $\mathrm{C}_{y,t}$. $\mathrm{X} \in \mathbb{R}^{d_x\times t}$ and $\mathrm{Y} \in \mathbb{R}^{d_y\times t}$ denote the matrices with t observed i.i.d. samples from the two views stacked as columns.

1.3. Problem Definition

The regularized version of Problem 1 can be formulated as follows. Given a pair of random vectors $\mathrm{x} \in \mathbb{R}^{d_x}$ and $\mathrm{y} \in \mathbb{R}^{d_y}$, distributed according to some unknown distribution $\mathscr{D}$, find the k-

1. www

dimensional subspaces where the projections of x and y are maximally correlated. With this notation, we can write the regularized CCA problem as follows.

$$\begin{aligned} \max_{\tilde{\mathrm{U}} \in \mathbb{R}^{d_x \times k}, \tilde{\mathrm{V}} \in \mathbb{R}^{d_y \times k}} \quad & \mathrm{Tr}\left(\tilde{\mathrm{U}}^\top \mathbb{E}[\mathrm{x}\mathrm{y}^\top] \tilde{\mathrm{V}}\right) \\ \text{subject to} \quad & \tilde{\mathrm{U}}^\top \mathbb{E}[\mathrm{x}\mathrm{x}^\top + r_x \mathrm{I}] \tilde{\mathrm{U}} = \mathrm{I}, \ \tilde{\mathrm{V}}^\top \mathbb{E}[\mathrm{y}\mathrm{y}^\top + r_y \mathrm{I}] \tilde{\mathrm{V}} = \mathrm{I} \end{aligned} \tag{2}$$

Since $\mathrm{C}_x^{\frac{1}{2}}$ and $\mathrm{C}_y^{\frac{1}{2}}$ are positive definite matrices one can make the simple change of variables $\mathrm{U} = \mathrm{C}_x^{\frac{1}{2}} \tilde{\mathrm{U}}$ and $\mathrm{V} = \mathrm{C}_y^{\frac{1}{2}} \tilde{\mathrm{V}}$ to get the equivalent problem

$$\begin{aligned} \max_{\mathrm{U} \in \mathbb{R}^{d_x \times k}, \mathrm{V} \in \mathbb{R}^{d_y \times k}} \quad & \mathrm{Tr}\left(\mathrm{U}^\top \mathbb{E}[\mathrm{x}\mathrm{x}^\top + r_x \mathrm{I}]^{-\frac{1}{2}} \, \mathbb{E}[\mathrm{x}\mathrm{y}^\top] \, \mathbb{E}[\mathrm{y}\mathrm{y}^\top + r_y \mathrm{I}]^{-\frac{1}{2}} \mathrm{V}\right) \\ \text{subject to} \quad & \mathrm{U}^\top \mathrm{U} = \mathrm{I}, \ \mathrm{V}^\top \mathrm{V} = \mathrm{I} \end{aligned} \tag{3}$$

Let $\Phi \in \mathbb{R}^{d_x \times k}$ and $\Psi \in \mathbb{R}^{d_y \times k}$ denote the top-k left and right singular vectors of $\mathrm{C}_x^{-\frac{1}{2}} \mathrm{C}_{xy} \mathrm{C}_y^{-\frac{1}{2}}$. Then, the optimum of Problem 2 is achieved at $\tilde{\mathrm{U}} = \mathrm{C}_x^{-\frac{1}{2}} \Phi$ and $\tilde{\mathrm{V}} = \mathrm{C}_y^{-\frac{1}{2}} \Psi$. An equivalent definition for the regularized empirical problem in the 1-dimensional setting is

$$\begin{aligned} \min_{\mathrm{u} \in \mathbb{R}^{d_x}, \mathrm{v} \in \mathbb{R}^{d_y}} \quad & \frac{1}{2t} \left\| \mathrm{u}^\top \mathrm{X} - \mathrm{v}^\top \mathrm{Y} \right\|_2^2 + r_x \left\| \mathrm{u} \right\|_2^2 + r_y \left\| \mathrm{v} \right\|_2^2 \\ \text{subject to} \quad & \mathrm{u}^\top \mathrm{C}_{x,t} \mathrm{u} = \mathrm{I}, \ \mathrm{v}^\top \mathrm{C}_{y,t} \mathrm{v} = \mathrm{I} \end{aligned} \tag{4}$$

This formulation motivates the alternating least squares (ALS) methods presented in Lu and Foster (2014); Ma et al. (2015); Wang et al. (2016).

1.4. Previous work

An approximate iterative least squares algorithm is proposed in (Lu and Foster, 2014). Their algorithm is based on alternating approximate projections of each view onto a k-dimensional subspace of the other using a least squares iteration. The paper does not provide any guarantee that the approximate solution will converge to the empirical risk minimizer. In (Ma et al., 2015), the authors propose AppGrad which is an alternating gradient descent procedure on Problem 4 followed by projection onto the constraint set. Their algorithm converges when initialized within a local neighborhood of the empirical estimates.

Two recent papers (Wang et al., 2016; Ge et al., 2016) propose very related approaches to solving the empirical problem in the batch setting and give global convergence guarantees for the resulting empirical risk minimization problem. Wang et al. (2016) consider an alternating least squares (ALS) procedure with appropriate normalization at each iteration. They leverage recent advances in designing fast algorithms for least squares regression based on accelerated first-order methods like SVRG (Johnson and Zhang, 2013). Ge et al. (2016) study CCA as a generalized eigenvalue problem. They consider a power iteration method which requires computing matrix inverse at each iteration. The update is reformulated as a convex optimization problem which is then solved approximately using an accelerated gradient descent method.

Finally, another recent paper by Allen-Zhu and Li (2016) proposes iterative algorithms for solving the empirical k-component CCA problem, which attain global convergence. The paper provides a first gap-free convergence result as well as improve on previous gap-dependent results. We would like to emphasize that none of the papers above give generalization bounds.

2. Proposed Method

Problem 3 is not a convex optimization problem – not only is the objective non-convex, the constraint set of orthogonal matrices is also non-convex. Here, we re-parametrize Problem 3 via the following variable substitute: $\mathrm{M} = \mathrm{U}\mathrm{V}^\top$. Furthermore, we take the convex hull of the constraint to give us the following convex relaxation:

$$\begin{aligned} \max_{\mathrm{M} \in \mathbb{R}^{d_x \times d_y}} \quad & \langle \mathrm{M}, \mathrm{C}_x^{-\frac{1}{2}} \mathrm{C}_{xy} \mathrm{C}_y^{-\frac{1}{2}} \rangle \\ \text{subject to} \quad & \|\mathrm{M}\|_2 \leq 1, \ \|\mathrm{M}\|_* \leq k \end{aligned} . \tag{5}$$

Similar relaxations have been previously considered to design stochastic approximation (SA) algorithms for principal component analysis (PCA) (Arora et al., 2013) and partial least squares (PLS) (Arora et al., 2016). These SA algorithms are instances of stochastic gradient descent – a popular choice for convex learning problems. Designing similar updates for the CCA problem is challenging since the gradient of the CCA objective (see Problem 5) w.r.t. M is $\mathrm{g} := \mathrm{C}_x^{-\frac{1}{2}} \mathrm{C}_{xy} \mathrm{C}_y^{-\frac{1}{2}}$, and it is not at all clear how one can design an unbiased estimator of the gradient g unless one knows the marginal distributions of x and y. If we knew the marginals, an unbiased estimator of the gradient would be $\mathrm{g}_t := \mathrm{C}_x^{-\frac{1}{2}} \mathrm{x}_t \mathrm{y}_t \mathrm{C}_y^{-\frac{1}{2}}$. We try to alleviate this problem by considering an "inexact" first order oracle

$$\partial_t := \mathrm{W}_{x,t} \mathrm{x}_t \mathrm{y}_t \mathrm{W}_{y,t} \approx \mathrm{g}_t,$$

which we show in Lemma 1 to converge in expectation to g_t, for the choice of $\mathrm{W}_{x,t} := \mathrm{C}_{x,t}^{-\frac{1}{2}}$ and $\mathrm{W}_{y,t} := \mathrm{C}_{y,t}^{-\frac{1}{2}}$ being the empirical estimates of the whitening transform matrices. We denote the true whitening transforms by $\mathrm{W}_x := \mathbb{E}[(\mathrm{x}\mathrm{x}^\top + r_x \mathrm{I})^{-\frac{1}{2}}]$ and $\mathrm{W}_y := \mathbb{E}[(\mathrm{y}\mathrm{y}^\top + r_y \mathrm{I})^{-\frac{1}{2}}]$.

We design stochastic approximation algorithms for solving Problem 5 where in each update we try to minimize the instantaneous loss measured on a single data point while trying to stay close to the previous iterate. In particular, at each iterate, we minimize the sum $\Delta(\mathrm{M}_t, \mathrm{M}) + \eta \mathbb{E}_{\mathscr{D}} \left[\ell \left(\mathrm{M}, \mathrm{g}_t \right) \right]$, over all feasible M, where $\Delta(\cdot, \cdot)$ is a divergence function between the successive iterates, ℓ is the instantaneous loss, and η is the learning-rate parameter controlling the tradeoff between loss and divergence. Often, the divergence function is defined in terms of a potential function. Here, we consider two different potential functions that lead to fundamentally different updates. Setting the potential function to be the Frobenius norm gives rise to divergence defined in terms of the Euclidean distance and yields additive updates (e.g. MSG in Section 2.1). Whereas, choosing the potential function to be von Neumann entropy gives rise to divergence function being quantum relative entropy resulting in multiplicative updates (e.g. MEG in Section 2.2). This is the standard approach to solving convex stochastic optimization problems with different variants subsumed in a common framework of stochastic mirror descent; see (Nedic and Lee, 2014; Nemirovskii et al., 1983; Tsuda et al., 2005) for more details. The focus here is to study the inexact versions of these algorithms for the CCA stochastic optimization problem which is non-convex.

2.1. Matrix Stochastic Gradient (MSG)

In this section, we consider stochastic mirror descent with the choice of potential function as the Frobenius norm. This yields an instance of matrix stochastic gradient (MSG) (Arora et al., 2013), except with an "inexact" gradient update:

$$\mathrm{M}_t = \mathscr{P}_F(\mathrm{M}_{t-1} + \eta_t \partial_t).$$

Algorithm 1 provides the pseudocode for the proposed inexact MSG for CCA. At each iteration of the algorithm, we update the empirical estimates of the covariance matrices which define the inexact gradient ∂_t. Lemma 1 guarantees that ∂_t (see steps 2-4 of Algorithm 1 for precise definition) converges in expectation to the unbiased estimator g_t. This is followed by projection onto the set of constraints in Problem 5 with respect to the Frobenius norm through the operator $\mathscr{P}_F(\cdot)$. We note that since the gradient ∂_t is a rank-1 matrix this update can be performed efficiently.

Algorithm 1 Matrix Stochastic Gradient (MSG)

Input: $\mathrm{M}_0, \{(\mathrm{x}_t, \mathrm{y}_t)\}_{t=1}^T, \eta$
Output: $\tilde{\mathrm{M}}$

1: **for** $t = 1, \cdots, T$ **do**
2: $\quad \mathrm{C}_{x,t} \leftarrow \frac{t-1}{t}\mathrm{C}_{x,t-1} + \frac{1}{t}\mathrm{x}_t\mathrm{x}_t^\top,\ \mathrm{W}_{x,t} \leftarrow \mathrm{C}_{x,t}^{-\frac{1}{2}}$
3: $\quad \mathrm{C}_{y,t} \leftarrow \frac{t-1}{t}\mathrm{C}_{y,t-1} + \frac{1}{t}\mathrm{y}_t\mathrm{y}_t^\top,\ \mathrm{W}_{y,t} \leftarrow \mathrm{C}_{y,t}^{-\frac{1}{2}}$
4: $\quad \partial_t \leftarrow \mathrm{W}_{x,t}\mathrm{x}_t\mathrm{y}_t\mathrm{W}_{y,t}$
5: $\quad \mathrm{M}_t \leftarrow \mathscr{P}_F(\mathrm{M}_t + \eta\partial_t)$ $\quad$ given in Arora et al. (2016)
6: **end for**
7: $\bar{\mathrm{M}} = \frac{1}{T}\sum_{t=1}^T \mathrm{M}_{t-1}$
8: $\tilde{\mathrm{M}} = \texttt{rounding}(\bar{\mathrm{M}})$

2.1.1. CONVERGENCE OF MSG FOR CCA

First, we show how to bound the error $\mathrm{E}_t = \mathrm{g}_t - \partial_t$ in spectral norm between the inexact gradient and g_t.

Lemma 1 *Assume that for all* $(\mathrm{x}_t, \mathrm{y}_t) \sim \mathscr{D}$ *we have* $\max\{\|\mathrm{x}_t\|^2, \|\mathrm{y}_t\|^2\} \leq B$. *There exists a constant* κ *independent of* t *such that for all* t *the following holds in expectation:*

$$\mathbb{E}_{\mathscr{D}}\left[\|\mathrm{E}_t\|_2\right] \leq \frac{\kappa}{\sqrt{t}}$$

Proof of lemma 1 is given in the Appendix. We are ready to state the main result of this section:

Theorem 2 After T iterations of MSG (Algorithm 1) with step size $\eta = \frac{2\sqrt{k}}{G\sqrt{T}}$, and starting at $\mathrm{M}^{(1)} = 0$,

$$\mathbb{E}[\langle \mathrm{M}_* - \tilde{\mathrm{M}}, \mathrm{C}_x^{-\frac{1}{2}}\mathrm{C}_{xy}\mathrm{C}_y^{-\frac{1}{2}}\rangle] \leq \frac{2\sqrt{k}G + 2k\kappa}{\sqrt{T}} \tag{6}$$

where κ is a universal constant given by lemma 1, M_* is the optimum of (5), expectation is with respect to the i.i.d. samples and rounding, and G is given by Lemma 16 such that $\|\partial_t\|_F \leq G$ for all the iterates $t = 1, \cdots, T$.

2.2. Matrix Exponentiated Gradient (MEG)

In this section, we consider matrix multiplicative weight updates for CCA. Multiplicative weights method is a generic algorithmic technique in which one updates a distribution over a set of interest by iteratively multiplying probability mass of elements (Kale, 2007). In our setting, the set is that

of d k-dimensional (paired) subspaces and the multiplicative algorithm is an instance of matrix exponentiated gradient (MEG) update. A motivation for considering MEG is the fact that for related problems, including principal component analysis (PCA) and partial least squares (PLS), MEG has been shown to yield fast optimistic[2] rates (Tsuda et al., 2005; Warmuth and Kuzmin, 2006a; Arora et al., 2016). We derive the MEG update as an instance of stochastic mirror descent for the choice of potential function being Quantum relative entropy.

Let $d := d_x + d_y$. Consider the self-adjoint dilation of the matrix $\mathbb{E}_{\mathscr{D}}[\mathrm{g}_t]$ (Tropp, 2012) $\mathrm{C} := \mathbb{E}_{\mathscr{D}}\left[\begin{pmatrix} 0 & \mathrm{g}_t \\ \mathrm{g}_t^\top & 0 \end{pmatrix}\right]$. If we assume that $\mathbb{E}_{\mathscr{D}}[\mathrm{g}_t]$ has no repeated singular values and its SVD is given by $\mathbb{E}_{\mathscr{D}}[\mathrm{g}_t] = \mathrm{U}\Sigma\mathrm{V}^\top$ then the eigen-decomposition of C is given by

$$\mathrm{C} = \frac{1}{2}\begin{pmatrix} \mathrm{U} & \mathrm{U} \\ \mathrm{V} & -\mathrm{V} \end{pmatrix}\begin{pmatrix} \Sigma & 0 \\ 0 & -\Sigma \end{pmatrix}\begin{pmatrix} \mathrm{U} & \mathrm{U} \\ \mathrm{V} & -\mathrm{V} \end{pmatrix}.$$

We can now formulate regularized CCA as

$$\begin{aligned} \max_{\mathrm{W}\in\mathbb{R}^{d\times k}} \quad & \mathrm{Tr}\left(\mathrm{W}\mathrm{W}^\top\mathrm{C}\right) \\ \text{subject to} \quad & \mathrm{W}^\top\mathrm{W} = \mathrm{I} \end{aligned}. \tag{7}$$

Problem 7 is not a convex optimization, but it admits the following convex relaxation by setting $\mathrm{M} = \frac{1}{k}\mathrm{W}\mathrm{W}^\top$:

$$\begin{aligned} \max_{\mathrm{M}\in\mathbb{R}^{d\times d}} \quad & \mathrm{Tr}\left(\mathrm{M}\mathrm{C}\right) \\ \text{subject to} \quad & \mathrm{M} \succeq 0, \|\mathrm{M}\|_2 \leq \frac{1}{k}, \mathrm{Tr}\left(\mathrm{M}\right) = 1 \end{aligned} \tag{8}$$

Following our discussion at the beginning of this section at each step of our algorithm we would like to update our current estimate of the projection matrix M_{t-1} by solving

$$\begin{aligned} \mathrm{M}_t := \operatorname*{argmin}_{\mathrm{M}\in\mathbb{R}^{d\times d}} \quad & \Delta\left(\mathrm{M}, \mathrm{M}_{t-1}\right) - \eta\mathrm{Tr}\left(\mathrm{M}\mathrm{C}_t\right) \\ \text{subject to} \quad & \mathrm{M} \succeq 0, \|\mathrm{M}\|_2 \leq \frac{1}{k}, \mathrm{Tr}\left(\mathrm{M}\right) = 1 \end{aligned} \tag{9}$$

where C_t is the self-adjoint dilation of g_t and $\Delta(x, y)$ is the quantum relative entropy between x and y. Setting the Lagrangian to 0 and solving we get the following update

$$\widehat{\mathrm{M}}_t = \frac{\exp\left(\log\left(\mathrm{M}_{t-1}\right) + \eta\mathrm{C}_t\right)}{\mathrm{Tr}\left(\exp\left(\log\left(\mathrm{M}_{t-1}\right) + \eta\mathrm{C}_t\right)\right)}, \quad \mathrm{M}_t = \mathscr{P}\left(\widehat{\mathrm{M}}_t\right) \tag{10}$$

where $\mathscr{P}$ denotes the projection onto the convex set of constraints in (9). In practice we replace the self-adjoint dilation of g_t by self-adjoint dilation of ∂_t which is denoted by $\tilde{\mathrm{C}}_t$.

2. If the optimal subspace has objective of the same size as ϵ, we get a convergence rate of $\mathcal{O}(\frac{1}{\epsilon})$ instead of $\mathcal{O}(\frac{1}{\epsilon^2})$.

Algorithm 2 MEG for CCA

Input: $\mathrm{M}_0, \{(\mathrm{x}_t, \mathrm{y}_t)\}_{t=1}^T, \eta$
Output: M
for $t = 1$ to T **do**
$\mathrm{C}_{x,t} \leftarrow \frac{t-1}{t}\mathrm{C}_{x,t-1} + \frac{1}{t}\mathrm{x}_t\mathrm{x}_t^\top,\ \mathrm{W}_{x,t} \leftarrow \mathrm{C}_{x,t}^{-\frac{1}{2}}$
$\mathrm{C}_{y,t} \leftarrow \frac{t-1}{t}\mathrm{C}_{y,t-1} + \frac{1}{t}\mathrm{y}_t\mathrm{y}_t^\top,\ \mathrm{W}_{y,t} \leftarrow \mathrm{C}_{y,t}^{-\frac{1}{2}}$
$\tilde{\mathrm{C}}_t \leftarrow \begin{pmatrix} 0 & \partial_t \\ \partial_t^\top & 0 \end{pmatrix} = \frac{1}{2}\begin{pmatrix} \mathrm{W}_{x,t}\mathrm{x}_t \\ \mathrm{W}_{y,t}\mathrm{y}_t \end{pmatrix}\begin{pmatrix} \mathrm{W}_{x,t}\mathrm{x}_t \\ \mathrm{W}_{y,t}\mathrm{y}_t \end{pmatrix}^\top - \frac{1}{2}\begin{pmatrix} \mathrm{W}_{x,t}\mathrm{x}_t \\ -\mathrm{W}_{y,t}\mathrm{y}_t \end{pmatrix}\begin{pmatrix} \mathrm{W}_{x,t}\mathrm{x}_t \\ -\mathrm{W}_{y,t}\mathrm{y}_t \end{pmatrix}^\top$
$\widehat{\mathrm{M}}_t \leftarrow \frac{\exp(\log(\mathrm{M}_{t-1}) + \eta\mathrm{C}_t)}{\mathrm{Tr}(\exp(\log(\mathrm{M}_{t-1}) + \eta\mathrm{C}_t))}$
$\mathrm{M}_t \leftarrow \mathscr{P}\left(\widehat{\mathrm{M}}_t\right)$ projection is given by algorithm 4 in Warmuth and Kuzmin (2008)
end for
$\bar{\mathrm{M}} = \frac{1}{T}\sum_{t=1}^T \mathrm{M}_{t-1}$
$\tilde{\mathrm{M}} = \texttt{rounding}\left(\bar{\mathrm{M}}\right)$

2.2.1. CONVERGENCE OF MEG FOR CCA

The analysis begins by proving a lemma similar to theorem 2 in Warmuth and Kuzmin (2006b).

Lemma 3 *After T iterations of algorithm 2 with a step size $\eta = \frac{1}{r}\log\left(1 + \sqrt{\frac{\log(d)}{LT}}\right)$ where $Tr\left(\mathrm{M}_*\tilde{\mathrm{C}}_t\right) \leq L$ for all t and $\mathrm{M}_0 = \frac{1}{d}\mathrm{I}$ we have*

$$\sum_{t=1}^T Tr\left(\mathrm{M}_*\tilde{\mathrm{C}}_t\right) - \sum_{t=1}^T Tr\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right) \leq 2\sqrt{rLT\log(d)} \tag{11}$$

where M_* is the optimal solution to (8) and r is chosen so that $\tilde{\mathrm{C}}_t - \lambda_{\min}\left(\tilde{\mathrm{C}}_t\right)\mathrm{I} \preceq r\mathrm{I}$.

Proof of lemma 3 is given in Appendix. Next we are going to give a bound on the quantity

$$\mathbb{E}_{\mathscr{D}}\left[\sum_{t=1}^T \frac{1}{T}\mathrm{Tr}\left(\mathrm{M}_*\mathrm{C}_t\right) - \sum_{t=1}^T \frac{1}{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\mathrm{C}_t\right)\right].$$

To do so, however, we need to give a bound in expectation of $\mathrm{Tr}\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right) - \mathrm{Tr}\left(\mathrm{M}^*\bar{\mathrm{E}}_t\right)$ where $\bar{\mathrm{E}}_t = \mathrm{C}_t - \tilde{\mathrm{C}}_t$.

Lemma 4 *With the same assumptions as in lemma 20 we have $\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[Tr\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right) - Tr\left(\mathrm{M}_*\bar{\mathrm{E}}_t\right)\right] \leq 2\frac{\kappa}{\sqrt{t}}$.*

We are finally ready to state and proof the main result of this section.

Theorem 5 *After T iterations of MEG Algorithm 2 starting from $\mathrm{M}_0 = \frac{1}{d}\mathrm{I}$ with step size $\eta = \frac{1}{r}\log\left(1 + \sqrt{\frac{\log(d)}{LT}}\right)$, we have that*

$$Tr\left(\mathrm{M}_*\mathrm{C}\right) - \mathbb{E}\left[Tr\left(\tilde{\mathrm{M}}\mathrm{C}\right)\right] \leq 2\sqrt{\frac{rL\log(d)}{T}} + 2\frac{\kappa}{\sqrt{T}}$$

where the expectation is taken with respect to the distribution and the internal randomization of the algorithm, $\mathrm{C} = \mathbb{E}_{\mathscr{D}}\left[\mathrm{C}_t\right]$, $\tilde{\mathrm{M}}$ *is the output of Algorithm 2,* r *is chosen such that* $0 \preceq \left(\mathrm{C}_t - \bar{\mathrm{E}}_t\right) - \lambda_{\min}\left(\left(\mathrm{C}_t - \bar{\mathrm{E}}_t\right)\right)\mathrm{I} \preceq r\mathrm{I}$, *and Tr* $\left(\mathrm{M}_* \left(\mathrm{C}_t - \bar{\mathrm{E}}_t\right)\right) \leq L$ *for all* $t = 1, \cdots, T$.

The proof of Theorem 5 is in the Appendix.

2.3. Computational Aspects

In this section we describe computational aspects and efficient implementation of both algorithms. The t-th iterations of both algorithms begins by updating the eigen-decomposition of $\mathrm{C}_{x,t-1}$ and $\mathrm{C}_{y,t-1}$. This can be done "efficiently" as follows. Let the ranks of $\mathrm{C}_{x,t-1} - r_x\mathrm{I}$ and $\mathrm{C}_{y,t-1} - r_y\mathrm{I}$ be $\mathrm{rk}_{x,t-1}$ and $\mathrm{rk}_{y,t-1}$. When we receive a new pair of points $(\mathrm{x}_t, \mathrm{y}_t)$ we first do a rank-1 update of the eigen-decomposition of $\mathrm{C}_{x,t-1} - r_x\mathrm{I}$ and $\mathrm{C}_{y,t-1} - r_y\mathrm{I}$ which takes respectively $\mathscr{O}\left(\mathrm{rk}^2_{x,t-1}d\right)$ and $\mathscr{O}\left(\mathrm{rk}^2_{y,t-1}d\right)$. We can then add back the regularization by $\mathrm{r}_x\mathrm{I}$ and $\mathrm{r}_y\mathrm{I}$. Having the updated eigen-decomposition of both auto-covariance operators it is now easy to compute $\mathrm{W}_{x,t}$ and $\mathrm{W}_{y,t}$ quickly by just inverting and taking the square roots of the eigenvalues of $\mathrm{C}_{x,t}$ and $\mathrm{C}_{y,t}$. These operations are done in $\mathscr{O}\left(\mathrm{rk}_{x,t}\right)$ and $\mathscr{O}\left(\mathrm{rk}_{y,t}\right)$ time respectively. Next we discuss the updates of the cross-covariance operators for MSG and MEG separately. For MSG the inexact stochastic gradient ∂_t is rank-1 so we can do a rank-1 SVD update on the estimate of the projection matrix. The projection $\mathscr{P}_F$ for MSG can also be efficiently done in $\mathscr{O}\left(d\log\left(d\right)\right)$ as discussed in Arora et al. (2016). For MEG we have the rank-2 self-adjoint dilation of ∂_t denoted by C_t, however, the update is a bit trickier as our initial projection matrix M_0 was set to a multiple of the identity which is a rank-d matrix. At every iteration one can keep the correct multiple of I which to subtract from M_{t-1} before doing the update making M_{t-1} a rank-rk_t matrix. Since M_t and $\log\left(\mathrm{M}_{t-1}\right) + \eta\mathrm{C}_t$ are symmetric matrices taking the logarithm and exponent amounts to just taking the logarithm and exponent of their eigenvalues. Also since C_t is a rank-2 matrix the update $\log\left(\mathrm{M}_{t-1}\right) + \eta\mathrm{C}_t$ takes $\mathscr{O}\left(d\mathrm{rk}^2_t\right)$ time. Finally the projection step $\mathrm{M}_t = \mathscr{P}\left(\widehat{\mathrm{M}}_t\right)$ takes $\mathscr{O}\left(d\right)$ time Warmuth and Kuzmin (2008). The last step of the algorithm $\mathrm{M} = \texttt{rounding}\left(\tilde{\mathrm{M}}\right)$ samples k eigen-values from the eigen-values making up the effective rank of M with probability proportional to their size which takes $\mathscr{O}\left(\mathrm{rk}_T\right)$

3. Generalization bounds for other CCA algorithms

We posed the CCA problem as maximizing the correlation objective. In this section, we try to better understand the computational tradeoffs offered by SAA by decomposing the "excess risk" – we therefore consider the equivalent problem of minimizing the negative of the correlation objective. In particular, we denote the true risk by $R(\mathrm{M}) = -\langle \mathrm{M}, \mathrm{C}_x^{-\frac{1}{2}}\mathrm{C}_{xy}\mathrm{C}_y^{-\frac{1}{2}}\rangle$ and the empirical risk based on t samples by $\widehat{R}(\mathrm{M}) = -\langle \mathrm{M}, \mathrm{C}_{x,t}^{-\frac{1}{2}}\mathrm{C}_{xy,t}\mathrm{C}_{y,t}^{-\frac{1}{2}}\rangle$. We consider the hypothesis class $\mathscr{M} := \{\mathrm{M} : \mathrm{Rank}\left(\mathrm{M}\right) = k,\ \|M\|_* = k,\ \|\mathrm{M}\|_1 = 1\}$. Let M_* and $\widehat{\mathrm{M}}_*$, respectively, be optimizers of R and $\widehat{R}$ in $\mathscr{M}$. Let the output of an approximation algorithm based on n samples be $\widehat{\mathrm{M}}_n$. We can then decompose the excess risk as follows:

$$\begin{aligned} R(\widehat{\mathrm{M}}_n) - R(\mathrm{M}_*) &= R(\widehat{\mathrm{M}}_n) - \widehat{R}(\widehat{\mathrm{M}}_n) + \widehat{R}(\widehat{\mathrm{M}}_n) - \widehat{R}(\mathrm{M}_*) + \widehat{R}(\mathrm{M}_*) - R(\mathrm{M}_*) \\ &\leq R(\widehat{\mathrm{M}}_n) - \widehat{R}(\widehat{\mathrm{M}}_n) + \widehat{R}(\widehat{\mathrm{M}}_n) - \widehat{R}(\widehat{\mathrm{M}}_*) + \widehat{R}(\mathrm{M}_*) - R(\mathrm{M}_*) \\ &\leq |R(\widehat{\mathrm{M}}_n) - \widehat{R}(\widehat{\mathrm{M}}_n)| + \widehat{R}(\widehat{\mathrm{M}}_n) - \widehat{R}(\widehat{\mathrm{M}}_*) + |\widehat{R}(\mathrm{M}_*) - R(\mathrm{M}_*)| \\ &\leq \widehat{R}(\widehat{\mathrm{M}}_n) - \widehat{R}(\widehat{\mathrm{M}}_*) + 2\max_{\mathrm{M}\in\mathscr{M}} |R(\mathrm{M}) - \widehat{R}(\mathrm{M})|, \end{aligned} \tag{12}$$

where the first inequality holds because $\widehat{\mathrm{M}}_*$ is the minimizer of $\widehat{R}$, in particular $\widehat{R}(\widehat{\mathrm{M}}_*) \leq \widehat{R}(\mathrm{M}_*)$. While MSG and MEG directly bound the excess risk, the previous works only bound optimization error $\widehat{R}(\widehat{\mathrm{M}}_*) - \widehat{R}(\widehat{\mathrm{M}}_n)$ as they are considering optimization of the empirical problem. In this section, we first provide uniform bounds on $\mathbb{E}_{\mathscr{D}}\left[|R(\mathrm{M}) - \widehat{R}(\mathrm{M})|\right]$. Subsequently, we give bounds on excess risk of methods presented in Wang et al. (2016), Ge et al. (2016) and Allen-Zhu and Li (2016).

In what follows, we let $\mathrm{X}_t \in \mathbb{R}^{d_x \times t}$ and $\mathrm{Y}_t \in \mathbb{R}^{d_y \times t}$ denote view-specific data matrices with data points stacked as columns, and $\mathrm{C}_{xy,t} = \frac{1}{t}\mathrm{X}_t\mathrm{Y}_t^\top$ gives the empirical cross-covariance matrix. Our results will depend crucially on spectral properties of the following symmetric dilation matrices:

$$\mathrm{T}_t := \begin{pmatrix} 0 & \mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t} \\ (\mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t})^\top & 0 \end{pmatrix}, \quad \mathrm{T} := \begin{pmatrix} 0 & \mathrm{W}_x\mathrm{C}_{xy}\mathrm{W}_y \\ (\mathrm{W}_x\mathrm{C}_{xy}\mathrm{W}_y)^\top & 0 \end{pmatrix}.$$

Without loss of generality, we assume that each of the methods works with regularized empirical auto-covariance matrices; Note that regularization has no effect on the convergence results and time complexity. First, we present a result that follows simply from Lemma 6 of Fukumizu et al. (2007).

Lemma 6 *Under the same assumptions as Lemma 1, we have that*

$$\mathbb{E}_{\mathscr{D}}\left[\left\|\frac{1}{t}\mathrm{W}_{x,t}\mathrm{X}_t\mathrm{Y}_t\mathrm{W}_{y,t} - \mathrm{W}_x\mathrm{C}_{xy}\mathrm{W}_y\right\|_2\right] \leq \frac{\tilde{\kappa}}{\sqrt{t}},$$

where $\tilde{\kappa}$ is a constant that does not depend on the number of samples t.

Furthermore, this gives a bound on the spectral norm of the difference of the self-adjoint dilations:

Corollary 7 *Under the same assumptions as Corollary 6, we have $\mathbb{E}_{\mathscr{D}}\left[\|\mathrm{T}_t - \mathrm{T}\|_2\right] \leq \frac{\tilde{\kappa}}{\sqrt{t}}$.*

The proof follows from relating the spectrum of a matrix with the spectrum of its self-adjoint dilation as discussed in section 2.2. Corollaries 6 and 7 yield the following uniform bound for all $\mathrm{M} \in \mathscr{M}$:

Corollary 8 $\max_{\mathrm{M}\in\mathscr{M}} \mathbb{E}_{\mathscr{D}}\left[|R(\mathrm{M}) - \widehat{R}(\mathrm{M})|\right] \leq \frac{k\tilde{\kappa}}{\sqrt{T}}$.

What follows are corollaries which give bounds on excess risk for Wang et al. (2016), Ge et al. (2016) and Allen-Zhu and Li (2016). In the bounds we provide we only show dependencies on the eigengap, combined dimensionality d and number of components. Note that we intentionally omit dependence on the conditional number as it is always bounded from above by $\frac{B}{\min\{r_x, r_y\}}$ and we consider it a constant. We begin with Wang et al. (2016) where we provide a bound for the top-1 setting as the algorithms in the paper are not specifically tailored for the top-k setting.

Corollary 9 *Using the ALS algorithm in Wang et al. (2016) with a least squares solver either AGD or SVRG we have $\mathbb{E}_{\mathscr{D}}\left[R(\widehat{\mathrm{M}}_t) - R(\mathrm{M}_*)\right] \leq \epsilon$ in time $\tilde{\mathcal{O}}\left(\frac{d}{\epsilon^2}\left(\frac{\sigma_1(\mathrm{T}_t)^2}{\sigma_1(\mathrm{T}_t)^2 - \sigma_2(\mathrm{T}_t)^2}\right)^2\right)$ where we hide logarithmic factors in ϵ.*

It is worth noting that we only consider the ALS algorithm with AGD and SVRG as a linear solver since ALS with ASVRG is only faster when $t < \frac{B}{\min\{r_x, r_y\}}$ (a setting which we do not consider) and the algorithms based on shift-and-invert preconditioning all give worse running time when considering generalization bounds. Next we give a generalization bound for Ge et al. (2016), however, since their results do not directly reflect what the relation between $\widehat{R}(\widehat{\mathrm{M}}_*)$ and $\widehat{R}(\widehat{\mathrm{M}}_t)$ is we need the following lemma first.

Lemma 10 *If* $\mathrm{u} \in \mathbb{R}^{d_x}$ *and* $\mathrm{v} \in \mathbb{R}^{d_y}$ *are the approximate j-th canonical directions for the empirical CCA problem on matrices* $\mathrm{C}_{x,t}, \mathrm{C}_{y,t}, \mathrm{C}_{xy,t}$ *returned by some algorithm and* $\mathrm{u}^* \in \mathbb{R}^{d_x}$ *and* $\mathrm{v}^* \in \mathbb{R}^{d_y}$ *are the exact j-th canonical directions for the empirical CCA problem and* $\mathrm{u}^\top \mathrm{C}_{x,t} u^* \geq 1 - \epsilon, \mathrm{v}^\top \mathrm{C}_{y,t} \mathrm{v}^* \geq 1 - \epsilon, \mathrm{u}^\top \mathrm{C}_{x,t} \mathrm{u} = 1, \mathrm{v}^\top \mathrm{C}_{y,t} \mathrm{v} = 1$ *then* $(\mathrm{u}^*)^\top \mathrm{C}_{xy,t} v^* - \mathrm{u}^\top \mathrm{C}_{xy,t} v \leq 2\epsilon B$.

We also need to introduce some notation given in Ge et al. (2016). First define the angle in the B-norm between two subspaces spanned by the columns of matrices W, V to be $\theta(\mathrm{W}, \mathrm{V}) = \arccos\left(\sigma_{\min}\left(\mathrm{V}^\top \mathrm{B} \mathrm{W}\right)\right)$. Now given two matrices W, V the authors say that W achieves an error of ϵ if $\mathrm{W}^\top \mathrm{B} \mathrm{W} = \mathrm{I}$ and $\sin(\theta(\mathrm{W}, \mathrm{V})) \leq \epsilon$. Denote the approximate solutions to the empirical CCA problem recovered by the algorithms in Ge et al. (2016) as $\tilde{\mathrm{U}} \in \mathbb{R}^{d_x \times k}, \tilde{\mathrm{V}} \in \mathbb{R}^{d_y \times k}$, s.t. $\tilde{\mathrm{U}}^\top \mathrm{C}_{x,t} \tilde{\mathrm{U}} = \mathrm{I}, \tilde{\mathrm{V}}^\top \mathrm{C}_{y,t} \tilde{\mathrm{V}} = \mathrm{I}$. If we denote the exact solutions by $\mathrm{U} \in \mathbb{R}^{d_x \times k}$ and $\mathrm{V} \in \mathbb{R}^{d_y \times k}$ then the authors claim that their algorithm achieves an error of ϵ if $\theta\left(\mathrm{U}, \tilde{\mathrm{U}}\right) \leq \epsilon$ in the $\mathrm{C}_{x,t}$ norm and $\theta\left(\mathrm{V}, \tilde{\mathrm{V}}\right) \leq \epsilon$ in the $\mathrm{C}_{y,t}$ norm.

Corollary 11 *Using the CCALin algorithm in Ge et al. (2016) with a least squares solver either AGD or SVRG we have* $\mathbb{E}_{\mathscr{D}}\left[R(\widehat{\mathrm{M}}_t) - R(\mathrm{M}_*)\right] \leq \epsilon$ *in time* $\tilde{\mathcal{O}}\left(\frac{kd}{\epsilon^2} \frac{|\sigma_k(\mathrm{T}_t)|}{|\sigma_k(\mathrm{T}_t) - \sigma_{k+1}(\mathrm{T}_t)|}\right)$ *where we hide logarithmic factors in* ϵ.

Finally we consider the work of Allen-Zhu and Li (2016). The results in their paper are given in a similar way to Ge et al. (2016), however, they consider the spectral norm between the k-dimensional approximate spaces spanned by the canonical directions and the complement of the k-dimensional exact spaces spanned by the canonical directions. It can be shown that this is equivalent to the way the results in Ge et al. (2016) are reported.

Corollary 12 *Using the LazyCCA algorithm in Allen-Zhu and Li (2016), the excess risk is bounded as* $\mathbb{E}_{\mathscr{D}}\left[R(\widehat{\mathrm{M}}_t) - R(\mathrm{M}_*)\right] \leq \epsilon$ *in time* $\tilde{\mathcal{O}}\left(\frac{kd}{\epsilon^2}\sqrt{\frac{|\sigma_k(\mathrm{T}_t)|}{|\sigma_k(\mathrm{T}_t) - \sigma_{k+1}(\mathrm{T}_t)|}}\right)$ *or* $\tilde{\mathcal{O}}\left(\frac{kd}{\epsilon^{\frac{5}{2}}}\right)$.

We omit the proof as it is similar to corollary 11.

4. Experiments

We provide experimental results for our proposed methods, MSG (Algorithm 1) and MEG (Algorithm 2), on a synthetic dataset and a real dataset, `Mediamill` (Snoek et al., 2006), consisting of paired observations of videos and corresponding commentary. We compare our algorithms against CCALin of Ge et al. (2016) and ALS CCA of Wang et al. (2016)[3]. All of the comparisons are given in terms of the CCA objective as a function of either CPU runtime or number of iterations[4]. Note that this puts MSG and MEG at a disadvantage when comparing in terms of the number of iterations as each iteration of CCALin and ALS CCA performs several passes over the data seen so far. The target dimensionality in our experiments is $k \in \{1, 2, 4, 8\}$. For regularization we fix $r_x = 0.01, r_y = 0.01$. For both MSG and MEG we set the step size at iteration t to be $\eta_t = \frac{0.1}{\sqrt{t}}$.

3. We run ALS only for $k = 1$ as the algorithm and the current implementation from the authors does not handle $k \geq 1$.
4. We consider the t-th iteration to occur when observing the t-th datapoint.

4.1. Synthetic dataset

Our synthetic dataset is constructed as follows. We sample a random vector $\mathrm{z} \in \mathbb{R}^{100}$ with the standard normal distribution, $\mathcal{N}(0, \mathrm{I})$. We generate a pair of observations $(\mathrm{x}_i, \mathrm{y}_i)$ for each sample z_i, $i = 1, \ldots, n$, using the following Gaussian noise model: $\mathrm{x} = \mathrm{U}^\top \mathrm{z} + \epsilon_x$ and $\mathrm{y} = \mathrm{V}^\top \mathrm{z} + \epsilon_y$, for fixed orthogonal matrices $\mathrm{U} \in \mathbb{R}^{100 \times d_x}$ and $\mathrm{V} \in \mathbb{R}^{100 \times d_y}$, and $\epsilon_x \sim \mathcal{N}\left(0, \sigma_x^2 \mathrm{I}\right)$, $\epsilon_y \sim \mathcal{N}\left(0, \sigma_y^2 \mathrm{I}\right)$ are random Gaussian noise with $\sigma_x = \sigma_y = 0.2$. We set $d_x = 100$, $d_y = 150$, and $n = 5000$. We note that canonical coefficients decay linearly in this setting which favors CCALin and ALS CCA. Experiments reveal that for $k = 1$, CCALin and ALS CCA exhibit worse performance when it comes to the per-iteration objective. ALS CCA has the worst performance in terms of CPU runtime and MEG, MSG, and CCALin have comparable performance. As k grows we see that the runtime performance drops significantly for CCALin compared to MSG and MEG, even for $k = 4$. As expected, MSG has worse per-iteration performance as k grows but in terms of CPU runtime it is the fastest. The plots are shown in Figure 1.

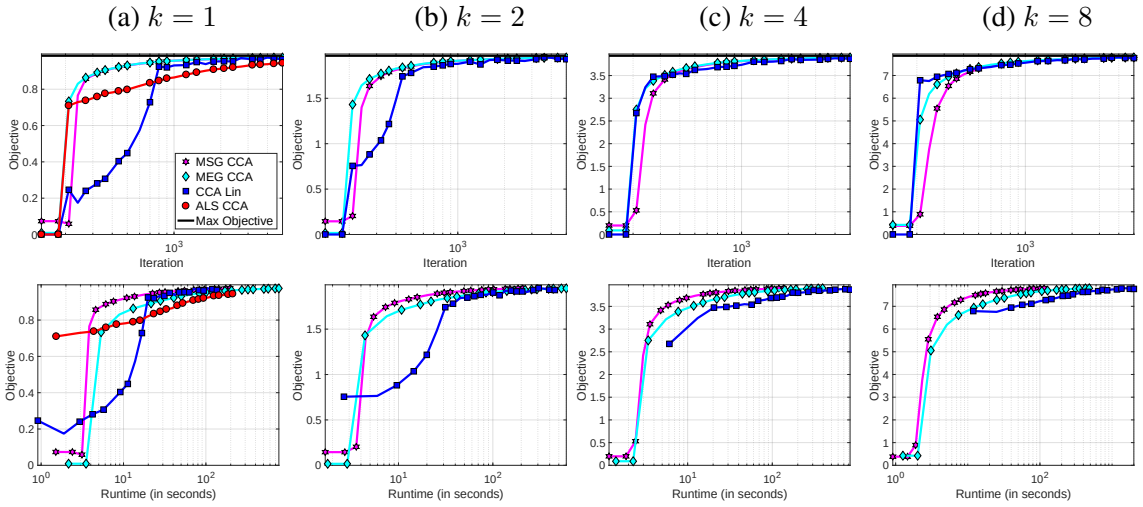

Figure 1: Comparisons of CCA-Lin, CCA-ALS, MSG, and MEG for CCA optimization on the synthetic dataset, in terms of the objective value as a function of iteration (top) and as a function of CPU runtime (bottom).

4.2. Mediamill

`Mediamill` is a multiview dataset consisting of $n = 10,000$ corresponding videos and text annotations with labels representing semantic concepts (Snoek et al., 2006). The image view consists of $d_x = 120$-dimensional visual features extracted from representative frames selected from videos, and the textual features are $d_y = 100$-dimensional. The estimated canonical coefficients have an exponential decay thereby favoring CCALin and ALS CCA as their convergence rates crucially depend on the eigengap. Therefore, as expected, both CCALin and ALS CCA outperform MEG and MSG when recovering the top CCA component, in terms of progress per-iteration and runtime. As k grows, we observe that the per-iteration progress of MEG starts matching the one of CCALin. For the CPU runtime comparison we notice that even when $k = 2$, MSG already matches the perfor-

mance of CCALin and as k grows we see that both MSG and MEG outperform CCALin. The plots are shown in Figure 2.

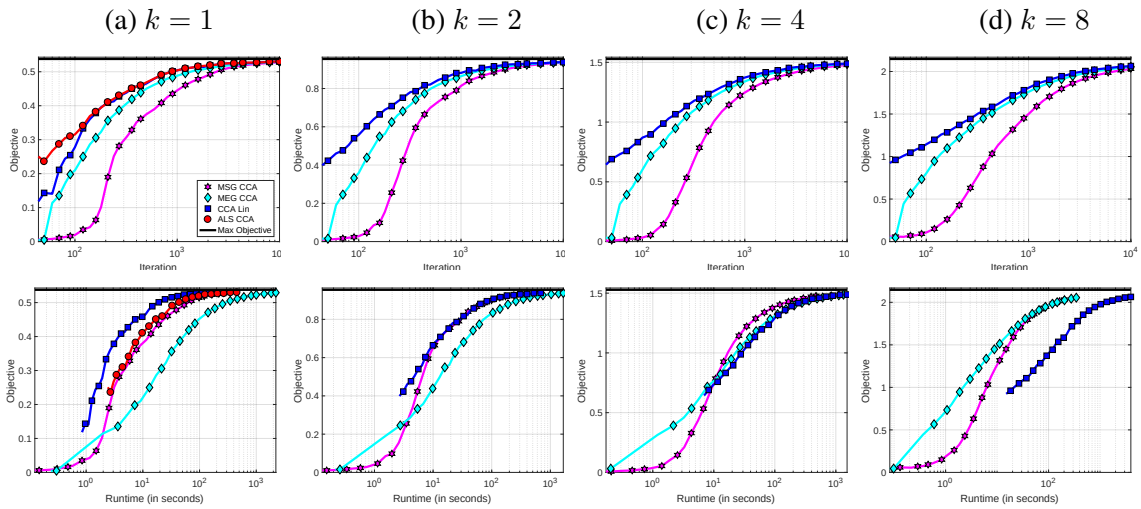

Figure 2: Comparisons of CCA-Lin, CCA-ALS, MSG, and MEG for CCA optimization on the MediaMill dataset, in terms of the objective value as a function of iteration (top) and as a function of CPU runtime (bottom).

5. Discussion

We study CCA as a stochastic optimization problem. We show that regularized CCA is efficiently PAC-learnable by providing analysis for two stochastic approximation algorithms. In particular, the proposed algorithms achieve ϵ-suboptimality in time $\mathcal{O}\left(\frac{1}{\epsilon^2}\right)$.

We provide generalization error bounds for existing state-of-the-art sample average approximation techniques. These bounds are either gap-dependent or are pessimistic[5].

Our work has implications to deep multiview representation learning, in particular deep variants of CCA (Andrew et al., 2013; Wang et al., 2015). These techniques have recently been applied successfully to numerous real-world problems. Stochastic approximation has been a key technology in the development of deep learning techniques that can leverage large volumes data; Unfortunately, deep CCA models, and its various extensions, depend on batch techniques such as LBFGS, which has been an impediment to their scalability to large datasets. The algorithms presented in this paper will pave the way to scalable deep learning techniques in the multiview setting.

References

Zeyuan Allen-Zhu and Yuanzhi Li. Doubly accelerated methods for faster cca and generalized eigendecomposition. *arXiv preprint arXiv:1607.06017*, 2016.

5. Allen-Zhu and Li (2016) gives a gap-free bound with sample complexity $\mathcal{O}\left(\frac{1}{\epsilon^{5/2}}\right)$.

Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1247–1255, 2013.

Raman Arora, Andrew Cotter, Karen Livescu, and Nathan Srebro. Stochastic optimization for pca and pls. In *Allerton Conference*, pages 861–868. Citeseer, 2012.

Raman Arora, Andy Cotter, and Nati Srebro. Stochastic optimization of pca with capped msg. In *Advances in Neural Information Processing Systems*, pages 1815–1823, 2013.

Raman Arora, Poorya Mianjy, and Teodor Marinov. Stochastic optimization for multiview representation learning using partial least squares. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1786–1794, 2016.

Sivaraman Balakrishnan, Kriti Puniyani, and John Lafferty. Sparse additive functional and kernel CCA. *International conference on machine learning (ICML)*, 2012.

Olivier Bousquet and Léon Bottou. The tradeoffs of large scale learning. In *Advances in neural information processing systems*, pages 161–168, 2008.

Shay B Cohen, Karl Stratos, Michael Collins, Dean P Foster, and Lyle Ungar. Spectral learning of latent-variable pcfgs. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 223–231. Association for Computational Linguistics, 2012.

Kenji Fukumizu, Francis R Bach, and Arthur Gretton. Statistical consistency of kernel canonical correlation analysis. *Journal of Machine Learning Research*, 8(Feb):361–383, 2007.

Rong Ge, Chi Jin, Sham M Kakade, Praneeth Netrapalli, and Aaron Sidford. Efficient algorithms for large-scale generalized eigenvector computation and canonical correlation analysis. *arXiv preprint arXiv:1604.03930*, 2016.

Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.

Satyen Kale. *Efficient algorithms using the multiplicative weights update method*. Princeton University, 2007.

Yichao Lu and Dean P Foster. Large scale canonical correlation analysis with iterative least squares. In *Advances in Neural Information Processing Systems*, pages 91–99, 2014.

Anastasia Lykou and Joe Whittaker. Sparse CCA using a lasso with positivity constraints. *Computational Statistics & Data Analysis*, 54(12):3144–3157, 2010.

Zhuang Ma, Yichao Lu, and Dean Foster. Finding linear structure in large datasets with scalable canonical correlation analysis. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 169–178, 2015.

Angelia Nedic and Soomin Lee. On stochastic subgradient mirror-descent algorithm with weighted averaging. *SIAM Journal on Optimization*, 24(1):84–107, 2014.

Arkadii Nemirovskii, David Borisovich Yudin, and Edgar Ronald Dawson. *Problem complexity and method efficiency in optimization*. Wiley, 1983.

Bernhard A Schmitt. Perturbation bounds for matrix square roots and pythagorean sums. *Linear algebra and its applications*, 174:215–227, 1992.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

Shai Shalev-Shwartz and Nathan Srebro. Svm optimization: inverse dependence on training set size. In *Proceedings of the 25th international conference on Machine learning*, pages 928–935. ACM, 2008.

Cees GM Snoek, Marcel Worring, Jan C Van Gemert, Jan-Mark Geusebroek, and Arnold WM Smeulders. The challenge problem for automated detection of 101 semantic concepts in multimedia. In *Proceedings of the 14th ACM international conference on Multimedia*, pages 421–430. ACM, 2006.

Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

Joel A Tropp. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.

Koji Tsuda, Gunnar Rätsch, and Manfred K Warmuth. Matrix exponentiated gradient updates for on-line learning and bregman projection. In *Journal of Machine Learning Research*, pages 995–1018, 2005.

Weiran Wang, Raman Arora, Karen Livescu, and Jeff A Bilmes. On deep multi-view representation learning. In *ICML*, pages 1083–1092, 2015.

Weiran Wang, Jialei Wang, Dan Garber, and Nati Srebro. Efficient globally convergent stochastic optimization for canonical correlation analysis. In *Advances in Neural Information Processing Systems*, pages 766–774, 2016.

Manfred K. Warmuth and Dima Kuzmin. Randomized PCA algorithms with regret bounds that are logarithmic in the dimension. In *NIPS'06*, 2006a.

Manfred K Warmuth and Dima Kuzmin. Online variance minimization. In *Learning theory*, pages 514–528. Springer, 2006b.

Manfred K Warmuth and Dima Kuzmin. Randomized online pca algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9(10), 2008.

Daniela M. Witten, Trevor Hastie, and Robert Tibshirani. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 2009.

Appendix A. MSG

First we introduce a structural result for perturbations of matrix square roots which is used in the proof of lemma 1.

Lemma 13 (Perturbation Bounds for Matrix Square Roots Schmitt (1992)) *Let* $\mathrm{A}_j \in \mathbb{R}^{n\times n}$ *with* $\mathrm{A}_j \succeq \mu_j^2 I$ *in the positive semi-definite order where* $j = 1, 2$. *Then* A_j *has a square root satisfying* $\mathrm{A}_j^{\frac{1}{2}} \succeq \mu_j I$ *and* $\|\mathrm{A}_1^{\frac{1}{2}} - \mathrm{A}_2^{\frac{1}{2}}\|_2 \leq \frac{1}{\mu_1+\mu_2}\|\mathrm{A}_1 - \mathrm{A}_2\|_2$

Proof [Proof of lemma 1] Let $\mathrm{A} = \mathrm{W}_x\mathrm{x}_t, \mathrm{B} = \mathrm{W}_y\mathrm{y}_t, \widehat{\mathrm{A}} = \mathrm{W}_{x,t}\mathrm{x}_t, \widehat{\mathrm{B}} = \mathrm{W}_{y,t}\mathrm{y}_t$. Notice that by regularization we have $\frac{1}{\sqrt{r}_x}\mathrm{I} \succeq \mathrm{W}_x, \frac{1}{\sqrt{r}_x}\mathrm{I} \succeq \mathrm{W}_{x,t}, \frac{1}{\sqrt{r}_y}\mathrm{I} \succeq \mathrm{W}_y, \frac{1}{\sqrt{r}_y}\mathrm{I} \succeq \mathrm{W}_{y,t}$.

$$\begin{aligned}
\mathbb{E}\left[\|\mathrm{E}_t\|_2\right] = \mathbb{E}\left[\|\mathrm{g}_t - \partial_t\|_2\right] &= \mathbb{E}\left[\left\|\mathrm{W}_x\mathrm{x}_t\mathrm{y}_t^\top\mathrm{W}_y - \mathrm{W}_{x,t}\mathrm{x}_t\mathrm{y}_t^\top\mathrm{W}_{y,t}\right\|_2\right] = \mathbb{E}\left[\left\|\mathrm{AB}^\top - \widehat{\mathrm{A}}\widehat{\mathrm{B}}^\top\right\|_2\right] \\
&= \mathbb{E}\left[\left\|\mathrm{AB}^\top - \mathrm{A}\widehat{\mathrm{B}}^\top + \mathrm{A}\widehat{\mathrm{B}}^\top - \widehat{\mathrm{A}}\widehat{\mathrm{B}}^\top\right\|_2\right] \\
&\leq \mathbb{E}\left[\|\mathrm{A}\|_2\left\|\mathrm{B} - \widehat{\mathrm{B}}\right\|_2 + \left\|\widehat{\mathrm{B}}\right\|_2\left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right] \qquad (13)
\end{aligned}$$

We first bound $\|\mathrm{A}\|_2$ and $\left\|\widehat{\mathrm{B}}\right\|_2$

$$\|\mathrm{A}\|_2 \leq \|\mathrm{W}_x\|_2\|\mathrm{x}_t\| \leq \sqrt{\frac{B}{r_x}}, \quad \left\|\widehat{\mathrm{B}}\right\|_2 \leq \|\mathrm{W}_{y,t}\|_2\|\mathrm{y}_t\| \leq \sqrt{\frac{B}{r_y}}$$

This implies that 13 is bounded by

$$\sqrt{\frac{B}{r_x}}\mathbb{E}\left[\left\|\mathrm{B} - \widehat{\mathrm{B}}\right\|_2\right] + \sqrt{\frac{B}{r_y}}\mathbb{E}\left[\left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right]. \qquad (14)$$

We now bound $\mathbb{E}\left[\left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right]$

$$\begin{aligned}
\mathbb{E}\left[\left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right] &\leq \mathbb{E}\left[\|\mathrm{W}_x - \mathrm{W}_{x,t}\|_2\|\mathrm{x}_t\|\right] \\
&\leq \sqrt{B}\mathbb{E}\left[\left\|\mathrm{C}_x^{-\frac{1}{2}} - \mathrm{C}_{x,t}^{-\frac{1}{2}}\right\|_2\right] \\
&= \sqrt{B}\mathbb{E}\left[\left\|\mathrm{C}_x^{-\frac{1}{2}}\left(\mathrm{C}_{x,t}^{\frac{1}{2}} - \mathrm{C}_x^{\frac{1}{2}}\right)\mathrm{C}_{x,t}^{-\frac{1}{2}}\right\|_2\right] \\
&\leq \sqrt{B}\mathbb{E}\left[\left\|\mathrm{C}_x^{-\frac{1}{2}}\right\|_2\left\|\mathrm{C}_{x,t}^{-\frac{1}{2}}\right\|_2\left\|\mathrm{C}_{x,t}^{\frac{1}{2}} - \mathrm{C}_x^{\frac{1}{2}}\right\|_2\right] \\
&\leq \frac{\sqrt{B}}{r_x}\mathbb{E}\left[\left\|\mathrm{C}_{x,t}^{\frac{1}{2}} - \mathrm{C}_x^{\frac{1}{2}}\right\|_2\right] \leq \frac{\sqrt{B}}{2r_x^{3/2}}\mathbb{E}\left[\|\mathrm{C}_{x,t} - \mathrm{C}_x\|_2\right] \qquad (15)
\end{aligned}$$

Where the last inequality follows from lemma 13. By lemma 14 $\mathbb{E}\left[\|\mathrm{C}_{x,t} - \mathrm{C}_x\|_2\right] \leq \sqrt{\frac{2B^2}{t}\log(d_x)} + \frac{2B}{3t}\log(d_x)$ and thus by equation 15

$$\begin{aligned}\mathbb{E}\left[\left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right] &\leq \frac{\sqrt{B}}{2r_x^{3/2}}\left\{\sqrt{\frac{2B^2}{t}\log(d_x)} + \frac{B}{3t}\log(d_x)\right\} \\ \mathbb{E}\left[\left\|\mathrm{B} - \widehat{\mathrm{B}}\right\|_2\right] &\leq \frac{\sqrt{B}}{2r_y^{3/2}}\left\{\sqrt{\frac{2B^2}{t}\log(d_y)} + \frac{B}{3t}\log(d_y)\right\}\end{aligned} \tag{16}$$

Finally 14 together with 16 implies that

$$\mathbb{E}\left[\|\mathrm{E}_t\|_2\right] \leq \frac{B^2}{2\sqrt{r_x r_y}}\left\{\frac{1}{r_y}\left\{\sqrt{\frac{2\log(d_y)}{t}} + \frac{3}{t}\log(d_y)\right\} + \frac{1}{r_x}\left\{\sqrt{\frac{2\log(d_x)}{t}} + \frac{3}{t}\log(d_x)\right\}\right\}$$

■

Lemma 14 *Under the assumptions of lemma 1* $\mathbb{E}_{\mathcal{D}}\left[\|\mathrm{C}_{x,t} - \mathrm{C}_x\|_2\right] \leq \sqrt{\frac{2B^2}{t}\log(d_x)} + \frac{B}{3t}\log(d_x)$.

Proof We bound the quantity by applying the Matrix Bernstein Inequality (Tropp (2015), Theorem 6.6.1). Set $\mathrm{X}_k = \frac{1}{t}\left(\mathrm{x}_k\mathrm{x}_k^\top - \mathrm{C}_{xx}\right)$. Where $\mathrm{C}_{xx} = \mathrm{C}_x - r_x\mathrm{I}$ so that $\mathbb{E}\left[\left\|\sum_{k=0}^{t-1}\mathrm{X}_k\right\|_2\right] = \mathbb{E}\left[\|\mathrm{C}_{x,t} - \mathrm{C}_x\|_2\right]$. To apply the inequality we need to verify that $\mathbb{E}\left[\mathrm{X}_k\right] = 0$ and $\|\mathrm{X}_k\|_2 \leq R$ and bound $\sigma^2 := \left\|\sum_k \mathbb{E}\left[X_k^2\right]\right\|_2$. Further, $\mathbb{E}\left[\mathrm{X}_k\right] = 0$ holds trivially. To bound $\|\mathrm{X}_k\|_2$

$$\|\mathrm{X}_k\|_2 = \frac{1}{t}\left\|\mathrm{x}_k\mathrm{x}_k^\top - \mathrm{C}_{xx}\right\|_2 \leq \frac{1}{t}\left(\left\|\mathrm{x}_k\mathrm{x}_k^\top\right\|_2 + \left\|\mathbb{E}\left[\mathrm{x}_k\mathrm{x}_k^\top\right]\right\|_2\right) \leq \frac{1}{t}\left(B + \mathbb{E}\left[\left\|\mathrm{x}_k\mathrm{x}_k^\top\right\|_2\right]\right) \leq \frac{2B}{t}$$

Finally we compute σ^2

$$\sum_k \mathbb{E}\left[\mathrm{X}_k^2\right] \preceq \frac{1}{t^2}\sum_k\left\{\mathbb{E}\left[B\mathrm{x}_k\mathrm{x}_k^\top\right] - \mathrm{C}_{xx}^2\right\} \preceq \frac{B}{t^2}\mathbb{E}\left[\sum_k \mathrm{x}_k\mathrm{x}_k^\top\right]$$

which implies $\sigma^2 \leq \frac{B^2}{t}$. By Matrix Bernstein's Inequality we have

$$\mathbb{E}\left[\left\|\sum_k \mathrm{X}_k\right\|_2\right] = \mathbb{E}\left[\|\mathrm{C}_{x,t} - \mathrm{C}_x\|_2\right] \leq \sqrt{\frac{2B^2}{t}\log(d_x)} + \frac{B}{3t}\log(d_x)$$

which completes the proof. ■

Lemma 15 $\sum_{t=1}^T \|\mathrm{E}_t\|_2 \leq 2\kappa\sqrt{T}$.

Proof We note that $\sum_{t=1}^T \frac{1}{\sqrt{t}} \leq \int_{t=1}^T \frac{1}{\sqrt{t}}\,dt + 1$. Substituting $z = \sqrt{t}$ and noting $dt = 2z\,dz$ we get

$$\int_{t=1}^T \frac{1}{\sqrt{t}}\,dt + 1 = \int_{z=1}^{\sqrt{T}} \frac{1}{z}\,2z\,dz + 1 = 2\sqrt{T} - 1 \leq 2\sqrt{T}.$$

■

Lemma 16 *With the same assumptions as in Lemma 1, we have* $\|\partial_t\|_F \leq B\sqrt{\frac{d_x d_y}{r_x r_y}}$.

Proof $\|\partial_t\|_F = \left\|\mathrm{W}_{x,t}\mathrm{x}_t\mathrm{y}_t^\top\mathrm{W}_{y,t}^\top\right\|_F \leq B\left\|\mathrm{W}_{x,t}\right\|_F\left\|\mathrm{W}_{y,t}\right\|_F \leq B\sqrt{\frac{d_x d_y}{r_x r_y}}$. ■

Proof [Proof of Theorem 2] We start the analysis by measuring the distance between the t-th iterate and the optimum, $D_t = \|\mathrm{M}_t - \mathrm{M}_*\|_F$.

$$\begin{aligned}
D_{t+1}^2 &= \|\mathrm{M}_{t+1} - \mathrm{M}_*\|_F^2 = \|\mathscr{P}_F(\mathrm{M}_t + \eta\partial_t) - \mathrm{M}_*\|_F^2 \\
&\leq \|\mathrm{M}_t + \eta\partial_t - \mathrm{M}_*\|_F^2 \\
&= \|\mathrm{M}_t - \mathrm{M}_*\|_F^2 + \eta^2\|\partial_t\|_F^2 + 2\eta\langle\mathrm{M}_t - \mathrm{M}_*, \mathrm{g}_t + \mathrm{E}_t\rangle \\
&\leq D_t^2 + \eta^2 G^2 + 2\eta\langle\mathrm{M}_t - \mathrm{M}_*, \mathrm{g}_t\rangle + 2\eta\langle\mathrm{M}_t - \mathrm{M}_*, \mathrm{E}_t\rangle \\
&\leq D_t^2 + \eta^2 G^2 + 2\eta\langle\mathrm{M}_t - \mathrm{M}_*, \mathrm{g}_t\rangle + 2\eta\|\mathrm{M}_t - \mathrm{M}_*\|_*\|\mathrm{E}_t\|_2 \\
&\leq D_t^2 + \eta^2 G^2 + 2\eta\langle\mathrm{M}_t - \mathrm{M}_*, \mathrm{g}_t\rangle + 4k\eta\|\mathrm{E}_t\|_2
\end{aligned}$$

Where the first inequality is due to projection onto a convex set in a Hilbert space being contractive and the third one is by Holder's inequality. The last inequality follows from $\|\mathrm{M}_t - \mathrm{M}_*\|_* \leq \|\mathrm{M}_t\|_* + \|\mathrm{M}_*\|_* \leq 2k$. Rearranging and dividing both sides by 2η we get

$$\langle\mathrm{M}_* - \mathrm{M}_t, \mathrm{g}_t\rangle \leq \frac{D_t^2 - D_{t+1}^2}{2\eta} + \frac{\eta}{2}G^2 + 2k\|\mathrm{E}_t\|_2$$

We average over T iterates, and note that $\sum_{t=1}^T D_t^2 - D_{t+1}^2 = D_1^2 - D_{T+1}^2 \leq D_1^2$, where the initial distance is bounded as follows:

$$D_1^2 = \|\mathrm{M}_1 - \mathrm{M}_*\|_F^2 = \|\mathrm{M}_1\|_F^2 + \|\mathrm{M}_*\|_F^2 - 2\langle\mathrm{M}_1, \mathrm{M}_*\rangle \leq k + k + 2\|\mathrm{M}_1\|_*\|\mathrm{M}_*\|_2 \leq 4k$$

taking the expectation from both sides, we get

$$\mathbb{E}[\langle\mathrm{M}_* - \tilde{\mathrm{M}}, \mathrm{C}_x^{-\frac{1}{2}}\mathrm{C}_{xy}\mathrm{C}_y^{-\frac{1}{2}}\rangle] \leq \frac{2k}{\eta T} + \frac{\eta G^2}{2} + \frac{2k\kappa\sqrt{T}}{T},$$

where we used lemma 15 to bound $\sum_{t=1}^T \|\mathrm{E}_t\|_2 \leq 2\kappa\sqrt{T}$, and choosing optimal learning rate $\eta = \frac{2\sqrt{k}}{G\sqrt{T}}$ we get the desired results. ■

Appendix B. MEG

Throughout our analysis we will need the following lemmas from Tsuda et al. (2005)

Lemma 17 *For any PSD matrix* A *and symmetric* B, C, $\mathrm{B} \preceq \mathrm{C}$ *implies* $Tr\left(\mathrm{AB}\right) \leq Tr\left(\mathrm{AC}\right)$

Lemma 18 *For any symmetric* A *such that* $0 \preceq \mathrm{A} \preceq \mathrm{I}$ *and any* $\rho_1, \rho_2 \in \mathbb{R}$ *the following holds*

$$\exp\left(\mathrm{A}\rho_1 + \left(\mathrm{I} - \mathrm{A}\right)\rho_2\right) \preceq \mathrm{A}\exp\left(\rho_1\right) + \left(\mathrm{I} - \mathrm{A}\right)\exp\left(\rho_2\right)$$

We also need the following lemma

Lemma 19 *For $x = 1 + \sqrt{\frac{R}{L}}$ the following holds*

$$\frac{-R + \log(x) L}{x - 1} \geq L - 2\sqrt{RL}$$

Proof This is a simple consequence of the fact that $\log(x) \geq (x-1) - (x-1)^2$ for $x \geq 1$ ■

Proof [Proof of Lemma 3] First we note that the update step 10 (after substituting C_t with $\tilde{\mathrm{C}}_t$) is invariant under perturbing the $\tilde{\mathrm{C}}_t$'s by a multiple of the identity Warmuth and Kuzmin (2006b) so we can assume that each $\tilde{\mathrm{C}}_t \succeq 0$. Also since $\max\left(\|x_t\|^2, \|y_t\|^2\right) \leq B$, $W_{x,t} \preceq \frac{1}{\sqrt{r_x}}\mathrm{I}$ and $W_{y,t} \preceq \frac{1}{\sqrt{r_y}}\mathrm{I}$ we can set $r = \frac{B}{\sqrt{r_x r_y}}$. We begin by considering the difference $\Delta(\mathrm{M}, \mathrm{M}_{t-1}) - \Delta\left(\mathrm{M}, \widehat{\mathrm{M}}_t\right)$ for any M satisfying the constraints 8.

$$\begin{aligned}
\Delta(\mathrm{M}, \mathrm{M}_{t-1}) - \Delta\left(\mathrm{M}, \widehat{\mathrm{M}}_t\right) &= \mathrm{Tr}\left(\mathrm{M}(\log(\mathrm{M}) - \log(\mathrm{M}_{t-1}))\right) - \mathrm{Tr}\left(\mathrm{M}\left(\log(\mathrm{M}) - \log\left(\widehat{\mathrm{M}}_t\right)\right)\right) \\
&= -\mathrm{Tr}\left(\mathrm{M}\left(\log(\mathrm{M}_{t-1}) - \log\left(\frac{\exp\left(\log(\mathrm{M}_{t-1}) + \eta\tilde{\mathrm{C}}_t\right)}{\mathrm{Tr}\left(\exp\left(\log(\mathrm{M}_{t-1}) + \eta\tilde{\mathrm{C}}_t\right)\right)}\right)\right)\right) \\
&= \eta\mathrm{Tr}\left(\mathrm{M}\tilde{\mathrm{C}}_t\right) - \log\left(\mathrm{Tr}\left(\exp\left(\log(\mathrm{M}_{t-1}) + \eta\tilde{\mathrm{C}}_t\right)\right)\right)
\end{aligned}$$

We now bound $\log\left(\mathrm{Tr}\left(\exp\left(\log(\mathrm{M}_{t-1}) + \eta\tilde{\mathrm{C}}_t\right)\right)\right)$. By Golden-Thompson's inequality we have

$$\mathrm{Tr}\left(\exp\left(\log(\mathrm{M}_{t-1}) + \eta\tilde{\mathrm{C}}_t\right)\right) \leq \mathrm{Tr}\left(\mathrm{M}_{t-1}\exp\left(\eta\tilde{\mathrm{C}}_t\right)\right)$$

Next we use lemma 18 on $\exp\left(\eta\tilde{\mathrm{C}}_t\right)$ with $\rho_0 = 0$ and $\rho_1 = r\eta$ since $0 \preceq \frac{\tilde{\mathrm{C}}_t}{r} \preceq \mathrm{I}$ to get $\exp\left(\eta\tilde{\mathrm{C}}_t\right) \preceq \frac{\tilde{\mathrm{C}}_t}{r}(\exp(r\eta) - 1) + \mathrm{I}$. By lemma 17 we now have

$$\mathrm{Tr}\left(\mathrm{M}_{t-1}\exp\left(\eta\tilde{\mathrm{C}}_t\right)\right) \leq \mathrm{Tr}\left(\mathrm{M}_{t-1} + \mathrm{M}_{t-1}\frac{\tilde{\mathrm{C}}_t}{r}(\exp(r\eta) - 1)\right)$$

which implies

$$\log\left(\mathrm{Tr}\left(\mathrm{M}_{t-1}\exp\left(\eta\tilde{\mathrm{C}}_t\right)\right)\right) \leq \log\left(1 + \frac{\mathrm{Tr}\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right)}{r}(\exp(r\eta) - 1)\right) \leq \frac{\mathrm{Tr}\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right)}{r}(\exp(r\eta) - 1)$$

where last inequality holds since $\log(1 + x) \leq x$. Thus

$$\Delta(\mathrm{M}, \mathrm{M}_{t-1}) - \Delta\left(\mathrm{M}, \widehat{\mathrm{M}}_t\right) \geq \eta\mathrm{Tr}\left(\mathrm{M}\tilde{\mathrm{C}}_t\right) - (\exp(r\eta) - 1)\frac{\mathrm{Tr}\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right)}{r}$$

equivalently

$$\mathrm{Tr}\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right) \geq r\frac{\Delta\left(\mathrm{M}, \widehat{\mathrm{M}}_t\right) - \Delta(\mathrm{M}, \mathrm{M}_{t-1}) + \eta\mathrm{Tr}\left(\mathrm{M}\tilde{\mathrm{C}}_t\right)}{\exp(r\eta) - 1}$$

By Generalized Pythagorean Theorem

$$\mathrm{Tr}\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right) \geq r\frac{\Delta\left(\mathrm{M},\mathrm{M}_t\right)-\Delta\left(\mathrm{M},\mathrm{M}_{t-1}\right)+\eta\mathrm{Tr}\left(\mathrm{M}\tilde{\mathrm{C}}_t\right)}{\exp\left(r\eta\right)-1}$$

Summing from $t=1$ to T and using the fact the Bregman divergence is positive we have

$$\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\tilde{\mathrm{C}}_t\right) \geq r\frac{-\Delta\left(\mathrm{M},\mathrm{M}_0\right)+\eta\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_*\tilde{\mathrm{C}}_t\right)}{\exp\left(r\eta\right)-1}$$

To complete the proof notice that $\Delta\left(\mathrm{M},\mathrm{M}_0\right)\leq\log\left(d\right)$ and apply lemma 19 with $\eta=\frac{1}{r}\log\left(1+\sqrt{\frac{r\log(d)}{LT}}\right)$. ■

Lemma 20 *Assuming* E_t *has no repeated singular values then* $-\frac{\kappa}{\sqrt{t}}\mathrm{I}\preceq\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]\preceq\frac{\kappa}{\sqrt{t}}\mathrm{I}$

Proof By the properties of self-adjoint dilation, we have $\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\left\|\mathrm{E}_t\right\|_2\right]=\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\left\|\bar{\mathrm{E}}_t\right\|_2\right]$. By Jensen's inequality and lemma 1 we have $\left\|\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]\right\|_2\leq\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\left\|\bar{\mathrm{E}}_t\right\|_2\right]=\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\left\|\mathrm{E}_t\right\|_2\right]\leq\frac{\kappa}{\sqrt{t}}$ and thus $-\frac{\kappa}{\sqrt{t}}\mathrm{I}\preceq\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]\preceq\frac{\kappa}{\sqrt{t}}\mathrm{I}$. ■

Lemma 21 *With the same assumptions as in lemma 20 we have*

$$\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[Tr\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right)-Tr\left(\mathrm{M}_*\bar{\mathrm{E}}_t\right)\right]\leq 2\frac{\kappa}{\sqrt{t}}$$

.

Proof Since M_{t-1} is independent of $(\mathrm{x}_t,\mathrm{y}_t)$ we have $\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\mathrm{Tr}\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right)\right]$ $=\mathrm{Tr}\left(\mathrm{M}_{t-1}\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]\right)$. From lemma 20 we know $\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]\preceq\frac{\kappa}{\sqrt{t}}\mathrm{I}$ and since $\mathrm{M}_{t-1}\succeq 0$ lemma 17 implies $\mathrm{Tr}\left(\mathrm{M}_{t-1}\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]\right)\leq\frac{\kappa}{\sqrt{t}}\mathrm{Tr}\left(\mathrm{M}_{t-1}\right)=\frac{\kappa}{\sqrt{t}}$. Similarly using that $-\frac{\kappa}{\sqrt{t}}\mathrm{I}\preceq\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\bar{\mathrm{E}}_t\right]$ we have $\mathbb{E}_{(\mathrm{x}_t,\mathrm{y}_t)}\left[\mathrm{Tr}\left(\mathrm{M}_*\bar{\mathrm{E}}_t\right)\right]\geq-\frac{\kappa}{\sqrt{t}}$ and the result follows. ■

Proof [Proof of Theorem 5] By lemma 3 we have

$$\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_*\left(\mathrm{C}_t-\bar{\mathrm{E}}_t\right)\right)-\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\left(\mathrm{C}_t-\bar{\mathrm{E}}_t\right)\right)\leq 2\sqrt{rLT\log\left(d\right)} \tag{17}$$

Let $\mathbb{E}_\tau\left[\cdot\right]$ denote the expectation w.r.t. $(\mathrm{x}_t,\mathrm{y}_t)_{t=1}^{\tau}$. We now compute the expectations of the two terms on the l.h.s. of 17

$$\mathbb{E}\left[\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_*\left(\mathrm{C}_t+\bar{\mathrm{E}}_t\right)\right)\right]=T\mathrm{Tr}\left(\mathrm{M}_*\mathrm{C}\right)-\sum_{t=1}^{T}\mathbb{E}\left[\mathrm{Tr}\left(\mathrm{M}_*\bar{\mathrm{E}}_t\right)\right] \tag{18}$$

The second term expands as follows

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\left(\mathrm{C}_t-\bar{\mathrm{E}}_t\right)\right)\right] &= \sum_{t=1}^{T}\mathrm{Tr}\left(\mathbb{E}_t\left[\mathrm{M}_{t-1}\left(\mathrm{C}_t-\bar{\mathrm{E}}_t\right)\right]\right)\\
&= \sum_{t=1}^{T}\mathrm{Tr}\left(\mathbb{E}_{t-1}\left[\mathbb{E}_t\left[\mathrm{M}_{t-1}\left(\mathrm{C}_t-\bar{\mathrm{E}}_t\right)|(\mathrm{x}_i,\mathrm{y}_i)_{i=1}^{t-1}\right]\right]\right)\\
&= \sum_{t=1}^{T}\mathrm{Tr}\left(\mathbb{E}_{t-1}\left[\mathrm{M}_{t-1}\right]\mathbb{E}_t\left[\left(\mathrm{C}_t-\bar{\mathrm{E}}_t\right)\right]\right) \qquad (19)\\
&= \mathbb{E}\left[\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\mathrm{C}\right)\right]-\sum_{t=1}^{T}\mathrm{Tr}\left(\mathbb{E}_{t-1}\left[\mathrm{M}_{t-1}\right]\mathbb{E}_t\left[\bar{\mathrm{E}}_t\right]\right)\\
&= \mathbb{E}\left[\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\mathrm{C}\right)\right]-\sum_{t=1}^{T}\mathbb{E}_{t-1}\left[\mathbb{E}_{\mathrm{x}_t,\mathrm{y}_t}\left[\mathrm{Tr}\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right)\right]\right]
\end{aligned}$$

Putting together 17, 18, 19 we have

$$\begin{aligned}
&T\mathrm{Tr}\left(\mathrm{M}_*\mathrm{C}\right)-\mathbb{E}\left[\sum_{t=1}^{T}\mathrm{Tr}\left(\mathrm{M}_{t-1}\mathrm{C}\right)\right]\\
&\leq 2\sqrt{rLT\log\left(d\right)}+\sum_{t=1}^{T}\left[\mathbb{E}_{t-1}\left[\mathbb{E}_{\mathrm{x}_t,\mathrm{y}_t}\left[\mathrm{Tr}\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right)\right]\right]-\mathbb{E}\left[\mathrm{Tr}\left(\mathrm{M}_*\bar{\mathrm{E}}_t\right)\right]\right] \qquad (20)\\
&= 2\sqrt{rLT\log\left(d\right)}+\sum_{t=1}^{T}\mathbb{E}_{t-1}\left[\mathbb{E}_{\mathrm{x}_t,\mathrm{y}_t}\left[\mathrm{Tr}\left(\mathrm{M}_{t-1}\bar{\mathrm{E}}_t\right)-\mathrm{Tr}\left(\mathrm{M}_*\bar{\mathrm{E}}_t\right)\right]\right]\\
&\leq 2\sqrt{rLT\log\left(d\right)}+\sum_{t=1}^{T}\frac{\kappa}{\sqrt{t}}\leq 2\sqrt{rLT\log\left(d\right)}+2\sqrt{T}\kappa
\end{aligned}$$

where the second inequality follows from lemma 21 and the last inequality follows from lemma 15. To finish the proof we divide both sides by T and take the expectation with respect to rounding. ■

Appendix C. Generalization Error Bounds

Proof [Proof of Lemma 6] This lemma simply follows from Lemma 6 of (Fukumizu et al., 2007). Here, we provide an alternative proof for the sake of academic discourse. We use the expectation symbol without a subscript whenever it is clear from the context. Let $\mathrm{A}=\frac{1}{\sqrt{t}}\mathrm{W}_x\mathrm{X}_t,\mathrm{B}=\frac{1}{\sqrt{t}}\mathrm{W}_y\mathrm{Y}_t,\widehat{\mathrm{A}}=\frac{1}{\sqrt{t}}\mathrm{W}_{x,t}\mathrm{X}_t,\mathrm{B}=\frac{1}{\sqrt{t}}\mathrm{W}_{y,t}\mathrm{Y}_t$. By regularization $\mathrm{W}_x\preceq\frac{1}{\sqrt{r_x}}\mathrm{I},\mathrm{W}_{x,t}\preceq\frac{1}{\sqrt{r_x}}\mathrm{I}$ and $\mathrm{W}_y\preceq\frac{1}{\sqrt{r_y}}\mathrm{I},\mathrm{W}_{y,t}\preceq\frac{1}{\sqrt{r_y}}\mathrm{I}$. We have

$$\begin{aligned}
&\mathbb{E}\left[\left\|\mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t}-\mathrm{W}_x\mathrm{C}_{xy}\mathrm{W}_y\right\|_2\right]\\
&=\mathbb{E}\left[\left\|\mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t}-\mathrm{W}_x\mathrm{C}_{xy,t}\mathrm{W}_y+\mathrm{W}_x\mathrm{C}_{xy,t}\mathrm{W}_y-\mathrm{W}_x\mathrm{C}_{xy}\mathrm{W}_y\right\|_2\right] \qquad (21)\\
&\leq\left\|\mathrm{W}_x\right\|_2\left\|\mathrm{W}_y\right\|_2\mathbb{E}\left[\left\|\mathrm{C}_{xy,t}-\mathrm{C}_{xy}\right\|_2\right]+\mathbb{E}\left[\left\|\mathrm{W}_x\mathrm{C}_{xy,t}\mathrm{W}_y-\mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t}\right\|_2\right]
\end{aligned}$$

We first bound $\|\mathrm{W}_x\|_2 \|\mathrm{W}_y\|_2 \mathbb{E}\left[\|\mathrm{C}_{xy,t} - \mathrm{C}_{xy}\|_2\right]$. Since $\mathrm{W}_x \preceq \frac{1}{\sqrt{r_x}}\mathrm{I}, \mathrm{W}_y \preceq \frac{1}{\sqrt{r_y}}\mathrm{I}$ we know it holds that $\|\mathrm{W}_x\|_2 \|\mathrm{W}_y\|_2 \leq \frac{1}{\sqrt{r_x r_y}}$. From matrix Bernstein's inequality Tropp (2015) theorem 6.1.1 $\mathbb{E}\left[\|\mathrm{C}_{xy,t} - \mathrm{C}_{xy}\|_2\right] \leq \sqrt{\frac{2B^2}{t}\log\left(d_x + d_y\right)} + \frac{B}{3t}\log\left(d_x + d_y\right)$ so

$$\|\mathrm{W}_x\|_2 \|\mathrm{W}_y\|_2 \mathbb{E}\left[\|\mathrm{C}_{xy,t} - \mathrm{C}_{xy}\|_2\right] \leq \frac{1}{\sqrt{r_x r_y}}\left(\sqrt{\frac{2B^2}{t}\log\left(d_x + d_y\right)} + \frac{B}{3t}\log\left(d_x + d_y\right)\right). \tag{22}$$

Next we bound $\mathbb{E}\left[\|\mathrm{W}_x \mathrm{C}_{xy,t} \mathrm{W}_y - \mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t}\|_2\right]$ exactly as in lemma 1

$$\begin{aligned}\mathbb{E}\left[\|\mathrm{W}_x \mathrm{C}_{xy,t} \mathrm{W}_y - \mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t}\|_2\right] &= \mathbb{E}\left[\left\|\mathrm{A}\mathrm{B}^\top - \widehat{\mathrm{A}}\widehat{\mathrm{B}}^\top\right\|_2\right] \\ &\leq \mathbb{E}\left[\|\mathrm{A}\|_2 \left\|\mathrm{B} - \widehat{\mathrm{B}}\right\|_2 + \left\|\widehat{\mathrm{B}}\right\|_2 \left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right]\end{aligned}$$

Since $\left\|\frac{1}{\sqrt{t}}\mathrm{X}_t\right\|_2 = \sqrt{\|\mathrm{C}_{x,t}\|_2} \leq \sqrt{B}$ we invoke the proof of lemma 1 to get that

$$\begin{aligned}\mathbb{E}\left[\|\mathrm{A}\|_2 \left\|\mathrm{B} - \widehat{\mathrm{B}}\right\|_2 + \left\|\widehat{\mathrm{B}}\right\|_2 \left\|\mathrm{A} - \widehat{\mathrm{A}}\right\|_2\right] \quad &\leq \frac{B^2}{2\sqrt{r_x r_y}}\left\{\frac{1}{r_y}\left\{\sqrt{\frac{2\log(d_y)}{t}} + \frac{3}{t}\log\left(d_y\right)\right\}\right. \\ &\left. + \frac{1}{r_x}\left\{\sqrt{\frac{2\log(d_x)}{t}} + \frac{3}{t}\log\left(d_x\right)\right\}\right\}\end{aligned} \tag{23}$$

Putting together 22 and 23 we get

$$\begin{aligned}\mathbb{E}\left[\left\|\frac{1}{t}\mathrm{W}_{x,t}\mathrm{X}_t\mathrm{Y}_t\mathrm{W}_{y,t} - \mathrm{W}_x\mathrm{C}_{xy}\mathrm{W}_y\right\|_2\right] \leq \frac{1}{\sqrt{r_x r_y}}&\left(\sqrt{\frac{2B^2}{t}\log\left(d_x + d_y\right)} + \frac{B}{3t}\log\left(d_x + d_y\right)\right) \\ &+ \frac{B^2}{2\sqrt{r_x r_y}}\left(\frac{1}{r_y}\left\{\sqrt{\frac{2\log\left(d_y\right)}{t}} + \frac{3}{t}\log\left(d_y\right)\right\}\right. \\ &\left. + \frac{1}{r_x}\left\{\sqrt{\frac{2\log\left(d_x\right)}{t}} + \frac{3}{t}\log\left(d_x\right)\right\}\right)\end{aligned}$$

■

Proof [Proof of Corollary 8] Simply follows from the fact that for any $\mathrm{M} \in \mathscr{M}$

$$\begin{aligned}\mathbb{E}_{\mathscr{D}}\left[\left|R(\mathrm{M}) - \widehat{R}(\mathrm{M})\right|\right] &= \mathbb{E}_{\mathscr{D}}\left[\left|\langle \mathrm{M}, \mathrm{C}_x^{-\frac{1}{2}}\mathrm{C}_{xy}\mathrm{C}_y^{-\frac{1}{2}} - \mathrm{C}_{x,t}^{-\frac{1}{2}}\mathrm{C}_{xy,t}\mathrm{C}_{y,t}^{-\frac{1}{2}}\rangle\right|\right] \\ &\leq \|\mathrm{M}\|_* \mathbb{E}_{\mathscr{D}}\left[\left\|\mathrm{C}_x^{-\frac{1}{2}}\mathrm{C}_{xy}\mathrm{C}_y^{-\frac{1}{2}} - \mathrm{C}_{x,t}^{-\frac{1}{2}}\mathrm{C}_{xy,t}\mathrm{C}_{y,t}^{-\frac{1}{2}}\right\|_2\right] \leq k \cdot \frac{\tilde{\kappa}}{\sqrt{T}}\end{aligned}$$

where the first inequality is due to Holder's and the second inequality follows from corollary 6. ■

Proof [Proof of Corollary 9] By theorem 2 in Wang et al. (2016) we know that for $k = 1$, $\widehat{R}(\widehat{\mathrm{M}}_t) \leq \widehat{R}(\widehat{\mathrm{M}}_*)(1 - 2\eta)$ in time $\tilde{\mathcal{O}}\left(dt\left(\frac{\sigma_1(\mathrm{T}_t)^2}{\sigma_1(\mathrm{T}_t)^2 - \sigma_2(\mathrm{T}_t)^2}\right)^2 \log\left(\frac{1}{\eta}\right)^2\right)$. Equivalently in the same

time we have $\widehat{R}(\widehat{M}_t) - \widehat{R}(\widehat{M}_*) \leq -2\eta\widehat{R}(\widehat{M}_*)$. Using the fact that $-\widehat{R}(\widehat{M}_*) \leq B$ and setting $\epsilon = 4B\eta$ we get $\widehat{R}(\widehat{M}_t) - \widehat{R}(\widehat{M}_*) \leq \frac{\epsilon}{2}$ and by corollary 8 the generalization error is bounded by $\frac{\epsilon}{2}$ in time $\mathcal{O}\left(\frac{k\tilde{\kappa}}{\epsilon^2}\right)$. ■

Proof [Proof of Lemma 10] The proof follows closely the one at the end of theorem 1 in Wang et al. (2016). Let $\phi = \mathrm{C}_{x,t}^{\frac{1}{2}}\mathrm{u}, \psi = \mathrm{C}_{y,t}^{\frac{1}{2}}\mathrm{v}$ and $\mathrm{T}_t := \mathrm{W}_{x,t}\mathrm{C}_{xy,t}\mathrm{W}_{y,t}$. Let also a_i and b_i denote the left and right singular vectors of T_t so that $\mathrm{a}_j = \mathrm{C}_{x,t}^{\frac{1}{2}}\mathrm{u}^*, \mathrm{b}_j = \mathrm{C}_{y,t}^{\frac{1}{2}}\mathrm{v}^*$. We have

$$\begin{aligned}
\mathrm{u}^\top \mathrm{C}_{xy,t} v = \phi^\top \mathrm{T}_t \psi &= \sigma_j\left(\mathrm{T}_t\right)\left(\phi^\top \mathrm{a}_j\right)\left(\psi^\top \mathrm{b}_j\right) + \sum_{i\neq j}\sigma_i\left(\mathrm{T}_t\right)\left(\phi^\top \mathrm{a}_i\right)\left(\psi^\top \mathrm{b}_i\right) \\
&\geq \sigma_j\left(\mathrm{T}_t\right)(1-\epsilon) - \sigma_1\left(\mathrm{T}_t\right)\sum_{i\neq j}|\phi^\top \mathrm{a}_i||\psi^\top \mathrm{b}_i| \\
&\geq \sigma_j\left(\mathrm{T}_t\right)(1-\epsilon) - \sigma_1\left(\mathrm{T}_t\right)\sqrt{\sum_{i\neq j}(\phi^\top \mathrm{a}_i)^2}\sqrt{\sum_{i\neq j}(\psi^\top \mathrm{b}_i)^2} \\
&\geq \sigma_j\left(\mathrm{T}_t\right)(1-\epsilon) - \sigma_1\left(\mathrm{T}_t\right)\epsilon \\
&\geq \sigma_j\left(\mathrm{T}_t\right) - B\epsilon - B\epsilon
\end{aligned} \tag{24}$$

Where the second inequality follows from Cauchy-Schwarz inequality, the third follows from the fact $\sum_{i=1}^{r}(\phi_i^\top \mathrm{a}_i)^2 \leq 1, \sum_{i=1}^{r}(\psi_i^\top \mathrm{b}_i)^2 \leq 1$ where r is the rank of T_t. ■

Proof [Proof of Corollary 11] By theorem 7 in Ge et al. (2016) we have that the proposed algorithm solves the empirical CCA problem on t points up to error ϵ as defined above in time $\tilde{\mathcal{O}}\left(tdk\frac{|\sigma_k(\mathrm{T}_t)|}{|\sigma_k(\mathrm{T}_t)-\sigma_{k+1}(\mathrm{T}_t)|}\right)$ (hiding logarithmic factors and factors involving the conditional number). Using the definition above this implies that in the same time we have $\sigma_{\min}\left(\tilde{\mathrm{U}}^\top \mathrm{C}_{x,t}\mathrm{U}\right) \geq \sqrt{1-\epsilon^2}$ and $\sigma_{\min}\left(\tilde{\mathrm{V}}^\top \mathrm{C}_{y,t}\mathrm{V}\right) \geq \sqrt{1-\epsilon^2}$. If we denote the i-th column of $\tilde{\mathrm{U}}, \mathrm{U}, \tilde{\mathrm{V}}, \mathrm{V}$ by respectively $\tilde{\mathrm{u}}_i, \mathrm{u}_i, \tilde{\mathrm{v}}_i, \mathrm{v}_i$ and use the fact $\sqrt{1-\epsilon^2} \geq 1-\epsilon$ for $\epsilon \leq 1$, we get that for $1 \leq i \leq k$, $\tilde{\mathrm{u}}_i^\top \mathrm{C}_{x,t}\mathrm{u}_i \geq 1-\epsilon$ and $\tilde{\mathrm{v}}_i^\top \mathrm{C}_{y,t}\mathrm{v}_i \geq 1-\epsilon$. By lemma 10 we have for $1 \leq i \leq k$, $u_i^\top \mathrm{C}_{xy,t} v_i - \tilde{u}_i^\top \mathrm{C}_{xy,t}\tilde{v}_i \leq 2\epsilon B$. Consequently we have that $\widehat{R}(\widehat{\mathrm{M}}_t) - \widehat{R}(\widehat{\mathrm{M}}_*) \leq 2k\epsilon B$. By corollary 8 the generalization error is bounded by $\frac{\epsilon}{2}$ in time $\mathcal{O}\left(\frac{k\tilde{\kappa}}{\epsilon^2}\right)$ and thus the claim of the corollary follows. ■